Rescaling Algorithms for Linear Programming Part I: Conic feasibility*

Daniel Dadush† dadush@cwi.nl

László A. Végh‡§ l.vegh@lse.ac.uk

Giacomo Zambelli ‡ g.zambelli@lse.ac.uk

arXiv:1611.06427v3 [math.OC] 12 Dec 2016

Abstract

We propose simple polynomial-time algorithms for two linear conic feasibility problems. For a matrix $A \in \mathbb{R}^{m\times n}$, the *kernel problem* requires a positive vector in the kernel of A, and the *image problem* requires a positive vector in the image of A^{T}. Both algorithms iterate between simple first order steps and rescaling steps. These rescalings steps improve natural geometric potentials in the domain and image spaces, respectively. If Goffin's condition measure $\hat{\rho}_A$ is negative, then the kernel problem is feasible and the worst-case complexity of the kernel algorithm is $O\left((m^3n + mn^2)\log|\hat{\rho}_A|^{-1}\right)$; if $\hat{\rho}_A > 0$, then the image problem is feasible and the image algorithm runs in time $O\left(m^2n^2\log\hat{\rho}_A^{-1}\right)$.

We also address the degenerate case $\hat{\rho}_A = 0$: we extend our algorithms for finding maximum support nonnegative vectors in the kernel of A and in the image of $A^{\top}$. We obtain the same running time bounds, with $\hat{\rho}_A$ replaced by appropriate condition numbers. In case the input matrix A has integer entries and total encoding length L, all algorithms are polynomial. Both full support and maximum support kernel algorithms run in time $O\left((m^3n + mn^2)L\right)$, whereas both image algorithms run in time $O\left(m^2n^2L\right)$. The standard linear programming feasibility problem can be easily reduced to either maximum support problems, yielding polynomial-time algorithms for Linear Programming.

1 Introduction

We consider two fundamental linear conic feasibility problems. Consider an $m \times n$ matrix A. In the *kernel problem* for A the goal is to find a positive vector in $\ker(A)$, whereas in the *image problem* the goal is to find a positive vector in $\mathrm{im}(A^{\mathsf{T}})$. These can be formulated by the following feasibility problems.

$$\begin{aligned} Ax &= 0 \\ x &> 0 \end{aligned} \qquad (K_{++}) \qquad\qquad A^{\mathsf{T}}y > 0 \qquad (I_{++})$$

We present simple polynomial-time algorithms for the kernel problem (K_{++}) and the image problem (I_{++}). Both algorithms combine a first order method with a geometric rescaling, which improve natural volumetric potentials.

The algorithms we propose fit into a line of research developed over the past 15 years [2, 3, 4, 8, 9, 10, 14, 21, 29, 30, 31, 37]. These are polynomial algorithms for Linear Programming that combine simple iterative updates, such as variants of perceptron [32] or of the relaxation method [1, 23], with some form of geometric rescaling. The current paper is the first in a paper series that gives new, more efficient algorithms based on this

*Partly based on the paper "Rescaled coordinate descent methods for Linear Programming" presented at the 18th Conference on Integer Programming and Combinatorial Optimization (IPCO 2016) [11].

†Centrum Wiskunde & Informatica. Supported by NWO Veni grant 639.071.510.

‡London School of Economics.

§Supported by EPSRC First Grant EP/M02797X/1.

technique and extends it to further applications. Our aim is to build a systematic theory of *geometric rescaling algorithms*.

Problems (K_{++}) and (I_{++}) have the following natural geometric interpretations. Let $a_1, \ldots, a_n$ denote the columns of the matrix A. A feasible solution to the (K_{++}) means that 0 is in the relative interior of the convex hull of the columns a_i, whereas a feasible solution to (I_{++}) gives a hyperplane that strictly separates 0 from the convex hull. For Goffin's condition measure $\hat{\rho}_A$, the value $|\hat{\rho}_A|$ is the distance of 0 from the relative boundary of the convex hull of the normalized vectors $a_i/\|a_i\|$. If $\hat{\rho}_A < 0$, then (K_{++}) is feasible, and if $\hat{\rho}_A > 0$, then (I_{++}) is feasible.

Both the kernel and image algorithms can be extended to the case $\hat{\rho}_A = 0$, that is, when 0 falls on the relative boundary of the convex hull of the a_i's. We address the following more general problems: to find a *maximum support* nonnegative vector in $\ker(A)$, and to find a maximum support nonnegative vector in $\mathrm{im}(A^{\mathsf{T}})$. Geometrically, these amount to identifying the face of the convex hull that contains 0 in its relative interior. By strong duality, the two maximum supports are complementary to each other.

To highlight the importance of the maximum support problems, let us note that an algorithm for either maximum support problem can be used directly for an LP feasibility problem of the form $Ax \leq b$ (i.e. general LP feasibility) via simple homogenization. While LP feasibility (and thus LP optimization) can also be reduced either to (K_{++}) or to (I_{++}) via standard perturbation methods (see for example [33]), this is not desirable for numerical stability. The maximum support problems provides fine-grained structural information on LP, and are crucial for exact LP algorithms (see e.g. [36]).

Previous work. We give a brief overview of geometric rescaling algorithms that combine first order iterations and rescalings. The first such algorithms were given by Betke [4] and by Dunagan and Vempala [14]. Both papers address the problem (I_{++}). The deterministic algorithm of Betke [4] combines a variant of Wolfe's algorithm with a rank-one update to the matrix A. Progress is measured by showing that the spherical volume of the cone $A^{\top}y \geq 0$ increases. This approach was further improved by Soheili and Peña [30], using different first order methods, in particular, a smoothed perceptron algorithm [25, 34]. Dunagan and Vempala [14] give a randomized algorithm, combining two different first order methods. They also use a rank-one update, but a different one from [4], and can show progress directly in terms of Goffin's condition measure $\hat{\rho}_A$. Extension of (I_{++}) to a conic setting were also given [3, 29]

For (K_{++}), as well as for the maximum support version, a rescaling algorithm was given by Chubanov [10], see also [21, 31]. A main iteration of the algorithm concludes that in the system $Ax = 0, 0 \leq x \leq \vec{e}$, one can identify at least one index i such that $x_i \leq \frac{1}{2}$ must hold for every solution. Hence the rescaling multiplies A from the right hand side by a diagonal matrix. (This is in contrast to the above mentioned algorithms, where rescaling multiplies the matrix A from the left hand side.) The first order iterations are von Neumann steps on the system defined by the projection matrix.

The algorithm [10] builds on previous work by Chubanov on binary integer programs and linear feasibility [8, 7], see also [2]. A more efficient variant of this algorithm was given in [37]. These algorithms use a similar rescaling, but for a non-homogeneous linear system, and the first order iterations are variants of the relaxation method.

Our contributions. We introduce new algorithms for (K_{++}) and (I_{++}) as well as for their maximum support versions. If $\hat{\rho}_A < 0$, that is, (K_{++}) is feasible, then the kernel algorithm runs in $O\left((m^3n + mn^2)\log|\hat{\rho}_A|^{-1}\right)$ arithmetic operations. If $\hat{\rho}_A > 0$, and thus (I_{++}) is feasible, and the image algorithm runs in $O\left(m^2n^2\log\hat{\rho}_A^{-1}\right)$ arithmetic operations. This can be improved to $O\left(m^3n\sqrt{\log n}\cdot\log\hat{\rho}_A^{-1}\right)$ using smoothing techniques [25, 34]. The rescalings improve volumetric potentials in the domain or in the image spaces, that act as proxies to the condition measure $|\hat{\rho}_A|$.

For the maximum support variants, we express the running time in terms of appropriate condition measures. These are close relatives of measures previously studied in the interior point literature. The kernel case uses measures θ_A, a variant of the symmetry measure introduced by Epelman and Freund [16], and ρ_A^*, which is closely related to Steward's χ_A [35]. The maximum support kernel algorithm runs in time $O((m^3n + mn^2)\log(\theta_A\rho_A^*)^{-1})$. In the image case, we introduce the measure $\hat{\omega}_A$, which is related to the dual condition measure of Vavasis and Ye [36]. The full support image algorithm runs in time $O\left(m^2n^2\log(n\hat{\omega}_A^{-1})\right)$. We also study a natural orthonormal preconditioning where many of the above condition measures become equivalent (see Section 4 and potentially improve.

Our algorithms use the *real model of computation* [5]. However, if the input matrix $A \in \mathbb{Z}^{m\times n}$ is integer of encoding length L, all condition numbers can be bounded by $2^{-O(L)}$, and therefore we obtain polynomial running bounds: $O\left((m^3n + mn^2)L\right)$ for both the full support and maximum support kernel algorithms, and $O\left(m^2n^2L\right)$ for both full support and maximum support image algorithms; or $O\left(m^3n\sqrt{\log n}\cdot L\right)$ using the smoothed variants.

Our algorithms improve on the best previous algorithms in the family of rescaling algorithms. For the maximum support kernel problem, the best previous algorithm was given by Chubanov [10], in time $O(n^4L)$. In applications, the number of rows m can be much smaller than the number of variables n. If $m < \sqrt{n}$, then our algorithm improves by a factor $\Omega(n^2/m)$. Further, we note that the condition numbers may provide much better running time estimates than the encoding size L used in [10]. For the image problem, our algorithm improves on Peña and Soheili's running-time [30] by a factor $\Omega(\sqrt{m\log(n)})$, and our smoothed variant improves by a factor $\Omega(n/\sqrt{m})$. A summary of running times is given in Table 1. For better comparability, we expressed all running times in terms of the encoding length L instead of using the various condition numbers.

Kernel problem	
Full support	Maximum support
$O(n^{18+3\varepsilon}\cdot L^{12+2\varepsilon})$ [7, 2]	
$O([n^5/\log n]\cdot L)$ [37]	
$O(n^4\cdot L)$ [10]	$O(n^4\cdot L)$ [10]
$O\left((m^3n + mn^2)\cdot L\right)$ this paper	$O((m^3n + mn^2)\cdot L)$ this paper

Image problem	
Full support	Maximum support
$O(m^3n^3\cdot L)$ [4]	
$O(m^4n\log m\cdot L)$[14]	
$O\left(m^{2.5}n^2\sqrt{\log n}\cdot L\right)$ [30]	
$O\left(m^3n\sqrt{\log n}\cdot L\right)$ this paper	$O\left(m^3n\sqrt{\log n}\cdot L\right)$ this paper

Table 1: Running time of geometric rescaling algorithms, in terms of encoding size L.

Our *kernel algorithm* can be seen as the most natural geometric rescaling algorithm for the kernel problem. We use a first order iteration used by Dunagan and Vempala [14], as well as the same rescaling. There are however substantial differences. Most importantly, Dunagan and Vempala assume $\hat{\rho}_A > 0$, and give a randomized algorithm to find a feasible solution to the image problem (I_{++}). In contrast, we assume $\hat{\rho}_A < 0$, and give a deterministic algorithm for finding a feasible solution the kernel problem (K_{++}), as well as for the more general maximum support problem. The key ingredient of our analysis is a new volumetric potential.

Our *image algorithm* is an improved version of Betke's [4] and Peña and Soheili's [30] algorithms. We introduce a new, more efficient rescaling, that enables to decrease the

number of first order iterations needed. In contrast to rank-1 updates used in all previous work, we use higher rank updates. For the maximum support image problem, we provide the first rescaling algorithm.

The full support kernel algorithm was first presented in the conference version [11]. The image algorithm and the maximum support variants for both the kernel and dual problems are new in this paper. The full support image algorithm was also independently obtained by Hoberg and Rothvoß [20].

While the current geometric rescaling methods are slower than current interior point or cutting plane methods, we believe there is much scope for improvement for future research. Also, they yield important insights into the structure of linear programs, for example, by identifying interesting subclasses of LPs that can be solved in strongly polynomial time (see for example [9]).

This paper is the first part in a series exploring geometric rescaling techniques for Linear Programming. In subsequent papers, we will provide a systematic study of the various rescaling algorithms and the relationships between them (highlighting the deep connections with the ellipsoid method), extensions of our algorithms to the oracle model, as well as applications to submodular function minimization and linear optimization.

The rest of the paper is structured as follows. Section 1.1 introduces notation and important concepts. Section 1.2 introduces the condition measures relevant to our algorithms, and Section 1.3 briefly surveys relevant first order methods. Section 2 presents the kernel algorithm in two variants. In Section 2.1 we give an algorithm for the full support problem (K_{++}) assuming it is feasible. This is extended in Section 2.2 to the maximum support case. Section 3 details the image algorithm and is subdivided in a similar manner. Section 4 discusses the relationship of our condition measures to others previously studied in the literature, and studies a natural preconditioning that leads to improved bounds.

1.1 Notation and preliminaries

For a natural number n, let $[n] = \{1, 2, \ldots, n\}$. For a subset $X \subseteq [n]$, we let $A_X \in \mathbb{R}^{m \times |X|}$ denote the submatrix formed by the columns of A indexed by X. For any non-zero vector $v \in \mathbb{R}^m$ we denote by $\hat{v}$ the normal vector in the direction of v, that is,

$$\hat{v} \stackrel{\text{def}}{=} \frac{v}{\|v\|}.$$

By convention, we also define $\hat{0} = 0$. We let $\hat{A} \stackrel{\text{def}}{=} [\hat{a}_1, \ldots, \hat{a}_n]$. Note that, given $v, w \in \mathbb{R}^m$, $\hat{v}^\mathsf{T}\hat{w}$ is the cosine of the angle between them.

Let $\mathbb{R}^n_+ = \{x \in \mathbb{R}^n : x \geq 0\}$ and $\mathbb{R}^n_{++} = \{x \in \mathbb{R}^n : x > 0\}$ denote the set of nonnegative and positive vectors in $\mathbb{R}^n$, respectively. For any set $H \subseteq \mathbb{R}^n$, we let $H_+ \stackrel{\text{def}}{=} H \cap \mathbb{R}^n_+$ and $H_{++} \stackrel{\text{def}}{=} H \cap \mathbb{R}^n_{++}$. These notations will be used in particular for the kernel and image spaces

$$\ker(A) \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \,:\, Ax = 0\}, \quad \text{im}(A^\mathsf{T}) \stackrel{\text{def}}{=} \{A^\mathsf{T} y \,:\, y \in \mathbb{R}^m\}.$$

Clearly, $\text{im}(A^\mathsf{T}) = \ker(A)^\perp$. Using this notation, (K_{++}) is the problem of finding a point in $\ker(A)_{++}$, and (I_{++}) amounts to finding a point in $\text{im}(A^\mathsf{T})_{++}$. By strong duality, (K_{++}) is feasible if and only if $\text{im}(A^\mathsf{T})_+ = \{0\}$, that is,

$$A^\mathsf{T} y \geq 0, \tag{I}$$

has no solution other than $y \in \ker(A^\mathsf{T})$. Similarly, (I_{++}) is feasible if and only if $\ker(A)_+ = \{0\}$, that is,

$$\begin{aligned} Ax &= 0 \\ x &\geq 0 \end{aligned} \tag{K}$$

has no solution other than $x = 0$. Let us define

$$\Sigma_A \stackrel{\text{def}}{=} \{y \in \mathbb{R}^m : A^\mathsf{T} y \geq 0\}$$

as the set of solutions to (I).

Given a vector $x \in \mathbb{R}^n$, the *support of* x is the subset of $[n]$ defined by $\text{supp}(x) \stackrel{\text{def}}{=} \{i \in [n] : x_i \neq 0\}$. Given any linear subspace H, we denote by $\text{supp}(H_+)$ the *maximum support* of H_+, that is, the unique inclusion-wise maximal element of the family $\{\text{supp}(x) : x \in H_+\}$. Note that, since H_+ is closed under summation, it follows that $\text{supp}(H_+) = \{i \in [n] : x_i > 0 \; \exists x \in H_+\}$.

Throughout the paper, we denote

$$S_A^* \stackrel{\text{def}}{=} \text{supp}(\ker(A)_+), \quad T_A^* \stackrel{\text{def}}{=} \text{supp}(\text{im}(A^\mathsf{T})_+). \tag{1}$$

When clear from the context, we will use the simpler notation S^* and T^*. Since $\ker(A)$ and $\text{im}(A^\mathsf{T})$ are orthogonal to each other, it is immediate that $S^* \cap T^* = \emptyset$. Furthermore, the strong duality theorem implies that $S^* \cup T^* = [n]$.

Let I_d denote the d dimensional identity matrix; we simply use I if the dimension is clear from the context. Let $\vec{e}_j$ denote the jth unit vector, and $\vec{e}$ denote the all-ones vector of appropriate dimension (depending on the context). For a point $p \in \mathbb{R}^d$ and $r > 0$, let $\mathbb{B}^d(p, r)$ denote the ball of radius r centered around p, where we write simply $\mathbb{B}(p, r)$ when the dimension is clear from the context. Let $\mathbb{B}^d = \mathbb{B}^d(0, 1)$ denote the d dimensional unit ball, and let $\nu_d = \text{vol}(\mathbb{B}^d)$ denote its volume.

Given any set C contained in $\mathbb{R}^d$, we denote by $\text{span}(C)$ the linear subspace of $\mathbb{R}^d$ spanned by the elements of C. If $C \subseteq \mathbb{R}^d$ has dimension r, we denote by $\text{vol}_r(C)$ the r-dimensional volume of C.

Projection matrices For any matrix $A \in \mathbb{R}^{m \times n}$, we denote by Π_A^I the *orthogonal projection matrix* to $\text{im}(A^\mathsf{T})$, and Π_A^K as the orthogonal projection matrix to $\ker(A)$. We recall that $\Pi_A^K + \Pi_A^I = I_n$, and they can be obtained as

$$\Pi_A^I = A^\mathsf{T}(AA^\mathsf{T})^+ A, \quad \Pi_A^K = I_n - A^\mathsf{T}(AA^\mathsf{T})^+ A,$$

where $(\cdot)^+$ denotes the Moore-Penrose pseudo-inverse. Note that, in order to compute Π_A^I and Π_A^K, one does not need to compute the pseudo-inverse of AA^T; instead, if we let B be a matrix comprised by $\text{rk}(A)$ many linearly independent rows of A, then $\Pi_A^I = B^\mathsf{T}(BB^\mathsf{T})^{-1}B$, which can be computed in $O(n^2 m)$ arithmetic operations.

Scalar products We will often need to use scalar products and norms other than the Euclidean ones. We use the following notation. We denote by $\mathbb{S}_+^d$ and $\mathbb{S}_{++}^d$ the sets of symmetric $d \times d$ positive semidefinite and positive definite matrices, respectively. Given $Q \in \mathbb{S}_{++}^d$, we denote by $Q^{\frac{1}{2}}$ the *square root* of Q, that is, the unique matrix in $\mathbb{S}_{++}^d$ such that $Q = Q^{\frac{1}{2}} Q^{\frac{1}{2}}$, and by $Q^{-\frac{1}{2}}$ the inverse of $Q^{\frac{1}{2}}$.

For $Q \in \mathbb{S}_{++}^d$ and two vectors $v, w \in \mathbb{R}^d$, we let

$$\langle v, w \rangle_Q \stackrel{\text{def}}{=} v^\top Q w, \quad \|v\|_Q \stackrel{\text{def}}{=} \sqrt{\langle v, v \rangle_Q}.$$

These define a scalar product and a norm over $\mathbb{R}^d$. We will use $\|\cdot\|_1$ for the $L1$-norm and $\|\cdot\|_2$ for the Euclidean norm. When there is no risk of confusion we will simply write $\|\cdot\|$ for $\|\cdot\|_2$. Further, for any $Q \in \mathbb{S}_{++}^d$, we define the ellipsoid

$$E(Q) \stackrel{\text{def}}{=} \{z \in \mathbb{R}^d : \|z\|_Q \leq 1\}.$$

For a linear subspace $H \subseteq \mathbb{R}^d$, we let $E_H(Q) = E(Q) \cap H$. Further, we define the projected determinant of Q on H as

$$\det_H(Q) \stackrel{\text{def}}{=} \det(W^\top Q W),$$

where W is a matrix whose columns form an orthonormal basis of H. Note that the definition is independent of the choice of the basis W. Indeed, if H has dimension r, then

$$\text{vol}_r(E_H(Q)) = \frac{\nu_r}{\sqrt{\det_H(Q)}}. \tag{2}$$

We conclude this section by summarizing well-known properties of matrix traces for later reference. We only prove parts (iii) and (iv).

Lemma 1.1.

(i) For any two matrices $X, Y \in \mathbb{R}^{k\times d}$, $\text{tr}(X^\top Y) = \text{tr}(XY^\top)$.

(ii) The trace is a linear function on square matrices.

(iii) For a positive semidefine matrix $X \in \mathbb{R}^{d\times d}$, $\det(I_d + X) \geq 1 + \text{tr}(X)$.

(iv) For any $X \in \mathbb{S}^d_{++}$, $\det(X)^{1/m} \leq \text{tr}(X)/m$.

Proof. **(iii):** Let $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_d \geq 0$ denote the eigenvalues of X, and let Λ denote the diagonal matrix with entries λ_i. Then

$$\det(I_d + X) = \prod_{i=1}^{d}(1 + \lambda_i) \geq 1 + \sum_{i=1}^{d} \lambda_i = 1 + \text{tr}(X),$$

where the equality is by the nonnegativity of the λ_i's.
(iv): Noting that $\det(X) = \prod_{i=1}^d \lambda_i$, this follows by the inequality of arithmetic and geometric means for the λ_i's. □ □

1.2 Condition measures

We will use several condition measures from the literature, and we also introduce natural new ones for our algorithms. We refer the reader to the survey [6] on condition measures and their role in the theory of interior point methods.

The most important measure in this context, and used in particular for the convergence analysis of the full support kernel algorithm, was first introduced by Goffin [18]. Given $A \in \mathbb{R}^{m\times n}$, we define

$$\rho_A \stackrel{\text{def}}{=} \max_{y \in \text{im}(A)\setminus\{0\}} \min_{j\in[n]} a_j^\mathsf{T} \hat{y}. \tag{3}$$

The Goffin [18] measure of A is the number

$$\hat{\rho}_A \stackrel{\text{def}}{=} \rho_{\hat{A}}.$$

Note that $\hat{\rho}_A$ depends only on the direction of the columns of A, not on their norm. Furthermore, we have $|\hat{\rho}_A|(\min_{j\in[n]} \|a_j\|) \leq |\rho_A| \leq |\hat{\rho}_A|(\max_{j\in[n]} \|a_j\|)$. We remark that, in the literature, A is typically assumed to have full row-rank (i.e. $\text{rk}(A) = m$), in which case y in the above definition ranges over all of $\mathbb{R}^m$. However, in several parts of the paper it will be convenient not to make such an assumption. The following Lemma summarizes well-known properties of ρ_A and $\hat{\rho}_A$; the proof will be given in the Appendix for completeness.

Lemma 1.2. *$|\rho_A|$ equals the distance of 0 from the relative boundary of* $\text{conv}(A)$*. Further,*

(i) $\rho_A < 0$ if and only if 0 is in the relative interior of $\text{conv}(A)$*, or equivalently, $S^* = [n]$.*

(ii) $\rho_A > 0$ *if and only if 0 is outside* $\mathrm{conv}(A)$*, or equivalently,* $T^* = [n]$*. In this case, the Goffin measure* $\hat{\rho}_A$ *equals the* width of the image cone Σ_A*, that is, the radius of the largest ball in* $\mathbb{R}^m$ *centered on the surface of the unit sphere and inscribed in* Σ_A*.*

Under the assumption that $\|a_i\| \leq 1$ for all $i \in [n]$, the full support kernel algorithm runs in time $O\left((m^3n + mn^2)\log(\rho_A^{-1})\right)$. If we initially renormalize to $\|a_i\| = 1$ for all $i \in [n]$, then ρ_A can be replaced by $\hat{\rho}_A$. The running time of the full support image algorithm can be bounded in terms of the Goffin measure, as $O\left(m^2n^2\log(\hat{\rho}_A^{-1})\right)$.

In the degenerate case, when 0 is on the boundary, that is $S^* \neq \emptyset$ and $T^* \neq \emptyset$, ρ_A will be 0. Therefore we need alternative condition measures for the maximum support kernel and image algorithms. We let

$$\theta_A \stackrel{\text{def}}{=} \min_{i \in S^*} \max\{\delta \,:\, -\delta a_i \in \mathrm{conv}(A)\}, \tag{4}$$

$$\rho_A^* \stackrel{\text{def}}{=} \min_{S \subseteq [n]} |\rho_{(A_S, -A_S)}|, \quad \text{and} \tag{5}$$

$$\hat{\omega}_A \stackrel{\text{def}}{=} \min_{i \in T^*} \max\left\{\hat{a}_i^\top y \,:\, y \in \mathbb{R}^m, A^\top y \geq 0, \|y\| \leq 1\right\}, \tag{6}$$

where $(A_S, -A_S)$ is the matrix obtained by juxtaposing the matrix A_S and its negative. The maximum support kernel algorithm runs in time $O\left((m^3n + mn^2)\log(\theta_A\rho_A^*)^{-1}\right)$, whereas the maximum support image algorithm runs in time $O\left(m^2n^2\log(n\hat{\omega}_A^{-1})\right)$; in the full support case, $\log(n\hat{\omega}_A^{-1})$ can be replaced by $\log\hat{\omega}_A^{-1}$. The measure $\hat{\omega}_A$ is related to the dual condition measure of Vavasis and Ye [36], see Claim 4.6. In the full support case, that is, if $T^* = [n]$, the measures $\hat{\omega}_A$ and $\hat{\rho}_A$ are essentially equivalent, as shown in the next statement.

Claim 1.3. *If* $\hat{\rho}_A > 0$*, then* $\hat{\rho}_A \leq \hat{\omega}_A \leq n\hat{\rho}_A$*.*

Proof. The inequality $\hat{\rho}_A \leq \hat{\omega}_A$ follows from the definition. For the other inequality, let

$$y^{(i)} := \arg\max\left\{\hat{a}_i^\top y \,:\, A^\top y \geq 0, \|y\| \leq 1\right\},$$

so that $\hat{\omega}_A = \min_{i\in[n]} \hat{a}_i^\top y^{(i)}$, and define $\bar{y} = \sum_{i=1}^n y^{(i)}$. Then $\hat{\rho}_A \geq \|\bar{y}\|^{-1} \min_{i\in[n]} \hat{a}_i^\mathsf{T}\bar{y} \geq \omega_A/n$, where the last inequality follows from the fact that $\|\bar{y}\| \leq n$ and $A^\mathsf{T}y^{(i)} \geq 0$ for all $i \in [n]$. □

Consequently, the running time of the full support image algorithm for $\hat{\rho}_A > 0$ can also be bounded as $O\left(m^2n^2\log(\hat{\rho}_A^{-1})\right)$, the running time stated in the Introduction.

In the kernel full support case (that is, $S^* = [n]$), the measure θ_A is the same as the symmetry measure of $\mathrm{conv}(A)$, introduced by Epelman and Freund [16]. If $S^* = [n]$, then $\theta_A = 1$ if and only if $\mathrm{conv}(A)$ is symmetric around the origin, and $\theta_A < 1$ otherwise. Regarding ρ_A^*, we will show in Lemma 4.5 that it is essentially equivalent to Stewart's measure χ_A [35].

We note that the maximum support algorithms require a priori knowledge of lower bounds on θ_A and $\hat{\omega}_A$ (but not on ρ_A^*); the running times are in terms of these lower bounds rather than the actual values. In contrast, the full support versions require no prior knowledge of ρ_A or $\hat{\rho}_A$. It remains an open problem to devise geometric rescaling algorithms for the maximum support problems that does not require a priori bounds on the condition numbers. Vavasis and Ye [36] gave such an interior point algorithm that solves both maximum support problems simultaneously.

All our algorithms are polynomial in the standard bit complexity model, assuming that we have an integer input matrix $A \in \mathbb{Z}^{m\times n}$. This follows by the following bounds; the standard proof is deferred to the Appendix.

Lemma 1.4. *Assume* $A \in \mathbb{Z}^{m\times n}$ *of total encoding length* L*. Then* $\theta_A, \rho_A^*, \hat{\omega}_A \geq 2^{-O(L)}$*. If* $\rho_A \neq 0$*, then* $|\rho_A|, |\hat{\rho}_A| \geq 2^{-O(L)}$*.*

Other measures have also been used implicitly in the literature on rescaling algorithms. For example, Chubanov's algorithm [10] for finding a maximum support solution to (K), is essentially about rescaling in order to increase the variant of σ_A^K defined with ∞-norm instead of 1-norm.

1.3 First order algorithms

Various first order methods are known for finding non-zero solutions to (K) or to (I). Most algorithms assume either the feasibility of (K_{++}) (that is, $S^* = [n]$ and $T^* = \emptyset$), or the feasibility of (I_{++}) (that is, $S^* = \emptyset$ and $T^* = [n]$). We outline the common update steps of these algorithms.

At every iteration, maintain a non-negative, non-zero vector $x \in \mathbb{R}^n$, and we let $y = Ax$. If $y = 0$, then x is a non-zero point in $\ker(A)_+$. If $A^\mathsf{T} y > 0$, then $A^\mathsf{T} y \in \mathrm{im}(A)_{++}$. Otherwise, choose an index $k \in [n]$ such that $a_k^\mathsf{T} y \leq 0$, and update x and y as follows:

$$y' := \alpha y + \beta \hat{a}_k; \quad x' := \alpha x + \frac{\beta}{\|a_k\|}\vec{e}_k, \tag{7}$$

where $\alpha, \beta > 0$ depend on the specific algorithm. Below we discuss various possible update choices.

von Neumann's algorithm We maintain at every iteration the condition that y is a convex combination of $\hat{a}_1, \ldots, \hat{a}_n$. The parameters $\alpha, \beta > 0$ are chosen so that $\alpha + \beta = 1$ and $\|y'\|$ is smallest possible. That is, y' is the point of minimum norm on the line segment joining y and $\hat{a}_k$. If we denote by y^t the vector at iteration t, and initialize $y^1 = \hat{a}_k$ for an arbitrary $k \in [n]$, a simple argument shows that $\|y^t\| \leq 1/\sqrt{t}$ (see Dantzig [13]). If 0 is contained in the interior of the convex hull, that is $\rho_A < 0$, Epelman and Freund [15] showed that $\|y^t\|$ decreases by a factor of $\sqrt{1 - \rho_A^2}$ in every iteration. Though the norm of y converges exponentially to 0, we note that this method may not actually terminate in finite time. If 0 is outside the convex hull however, that is, $\rho_A > 0$, then the algorithm terminates after at most $1/\rho_A^2$ iterations. A recent result by Peña, Soheili, and Rodriguez [28] gives a variant of the algorithm with a provable convergence guarantee in the case $\rho_A = 0$, that is, if 0 is on the boundary of the convex hull.

Betke [4] gave a polynomial time algorithm, based on a combinatorial variant of von Neumann's update, for the case $T^* = [n]$. Chubanov uses von Neumann's update on the columns of the projection matrix Π_A^K, and is able to solve the maximum support problem in time $O(n^4 L)$.

Note that von Neumann's algorithm is the same as the Frank-Wolfe conditional gradient descent method [17] with optimal step size for the quadratic program $\min \|Ax\|^2$ s.t. $\vec{e}^\mathsf{T} x = 1, x \geq 0$.

Perceptron algorithm The perceptron algorithm chooses $\alpha = \beta = 1$ at every iteration. If $\rho_A > 0$, then, similarly to the von Neumann algorithm, the perceptron algorithm terminates with a solution to the system (I_{++}) after at most $1/\rho_A^2$ iterations (see Novikoff [26]). The perceptron and von Neumann algorithms can be interpreted as duals of each other, see Li and Terlaky [22].

Peña and Soheili gave a smoothed variant of the perceptron update which guarantees termination in time $O(\sqrt{\log n}/\rho_A)$ iterations [34], and showed how this gives rise to a polynomial-time algorithm [30] using the rescaling introduced by Betke in [4]. The same iteration bound $O(\sqrt{\log n}/\rho_A)$ was achieved by Yu et al. [39] by adapting the Mirror-Prox algorithm of Nemirovski [24].

With a normalization to $e^\mathsf{T} x = 1$, perceptron implements the Frank-Wolfe algorithm for the same system $\min \|Ax\|^2$ s.t. $\vec{e}^\mathsf{T} x = 1, x \geq 0$, with step length $1/(k+1)$ at iteration

k. An alternative view is to interpret perceptron as a coordinate descent method for the system $\min \|Ax\|^2$ s.t. $x \geq \vec{e}$, with fixed step length 1 at every iteration.

Dunagan-Vempala algorithm The first order method used in [14] chooses $\alpha = 1$ and $\beta = -(\hat{a}_k^\mathsf{T} y)$. The choice of β is the one that makes $\|y'\|$ the smallest possible when $\alpha = 1$. It can be readily computed that

$$\|y'\| = \|y\|\sqrt{1 - (\hat{a}_k^\mathsf{T} \hat{y})^2}. \tag{8}$$

In particular, the norm of y' decreases at every iteration, and the larger is the angle between a_k and y, the larger the decrease. If $\rho_A < 0$, then $|\hat{a}_k^\mathsf{T} \hat{y}| \geq |\rho_A|$, therefore this guarantees a decrease in the norm of at least $\sqrt{1 - \rho_A^2}$.

This is a coordinate descent for the system $\min \|Ax\|^2$ s.t. $x \geq \vec{e}$, but with the optimal step length. One can also interpret it as the Frank-Wolfe algorithm with the optimal step length for the same system.[1]

Our kernel algorithm uses Dunagan-Vempala updates, and the image algorithm uses von Neumann updates. We combine the Dunagan-Vempala updates with a simple projection step, that guarantees finite convergence for finding a solution to (K_{++}) if $\rho_A < 0$, even without rescaling. This is a relevant contribution in the context of first order algorithms, since von Neumann's algorithm does not have finite convergence; this aspect is discussed by Li and Terlaky [22]. Dantzig [12] proposed a finitely converging variant of von Neumann's algorithm, but this involves running the algorithm $m + 1$ times, and also an explicit lower bound on the parameter $|\rho_A|$. Our algorithm does not incur a running time increase compared to the original variant, and does not require such a bound.

It is interesting to note that Dunagan and Vempala [14] use these updates for a different purpose. They address the opposite setting $\rho_A > 0$, and use two first order methods as subroutines: Perceptron for finding a solution once ρ_A is large, and these updates steps for the purpose of finding a good rescaling direction.

2 The Kernel Algorithm

Section 2.1 presents the full support kernel algorithm, that is, solving (K_{++}). This is followed by the maximum support kernel algorithm in Section 2.2. Besides addressing different problem settings and using different condition numbers, the technical frameworks are also different. In the full support kernel algorithm, rescalings are applied directly to the matrix A. In the maximum support variant, we keep the original matrix throughout, but modify the scalar product at every rescaling. These two frameworks are equivalent, and the algorithms can be easily translated from one to the other. Rescaling the matrix directly provides a very clear geometric intuition for the simpler full support setting. On the other hand, modifying the scalar product is more natural for analyzing the maximum support problem. We note that both the full support and maximum support versions of the image algorithm in Section 3 use the scalar product modification framework.

Both versions of the kernel algorithm maintain a vector $x \geq \vec{e}$ and the conic combination $y = Ax$. At every iteration, the algorithm chooses one of two steps: a Dunagan-Vempala (DV) update, if this results in a "substantial" decrease in the norm of the current vector $y = Ax$, or a geometric rescaling aimed at improving the condition measure. We use the volume of the polytope P_A defined by

$$P_A \stackrel{\text{def}}{=} \text{conv}(\hat{A}) \cap (-\text{conv}(\hat{A})). \tag{9}$$

[1]The Frank-Wolfe method is originally described for a compact set, but the set here is unbounded. Nevertheless, one can easily modify the method by moving along an unbounded recession direction.

The volume of P_A will be used as a proxy to $|\hat{\rho}_A|$ in the full support version. Recall from Lemma 1.2 that $|\hat{\rho}_A|$ is the radius of the largest ball around the origin inscribed in $\mathrm{conv}(\hat{A})$. This ball must be contained in P_A. Below we state a lemma that shows the relationship between P_A and S^*. The easy proof is postponed to the Appendix.

Lemma 2.1. *Let $A \in \mathbb{R}^{m\times n}$ and $S^* = S^*_A$. Then $\mathrm{span}(P_A) = \mathrm{im}(A_{S^*})$, and $P_A = P_{A_{S^*}}$.*

Throughout the kernel and image algorithms, we use the parameter

$$\varepsilon \stackrel{\text{def}}{=} \frac{1}{11m}. \tag{10}$$

2.1 Full support case

Algorithm 1 Full Support Kernel Algorithm

Input: A matrix $A \in \mathbb{R}^{m\times n}$ such that $\rho_A < 0$ and $\|a_j\| \le 1$ for all $j \in [n]$.
Output: A solution to the system (K_{++}).

Compute $\Pi := \Pi^K_A = I_n - A^\mathsf{T}(AA^\mathsf{T})^+A$.
Set $x_j := 1$ for all $j \in [n]$, and $y := Ax$.
while $\Pi x \not> 0$ **do**
    Let $k := \arg\min_{j\in[n]} \hat{a}_j^\mathsf{T}\hat{y}$;
    if $\hat{a}_k^\mathsf{T}\hat{y} < -\varepsilon$ **then**
        update $x := x - \dfrac{a_k^\mathsf{T}y}{\|a_k\|^2}\vec{e}_k$; $\quad y := y - (\hat{a}_k^\mathsf{T}y)\hat{a}_k$;
    else
        rescale $A := \left(I_m + \hat{y}\hat{y}^\mathsf{T}\right)A$; $\quad y := 2y$;
return $\bar{x} = \Pi x$ as a feasible solution to (K_{++}).

The Full Support Kernel Algorithm (Algorithm 1) solves the full support problem (K_{++}), that is, it finds a point in $\ker(A)_{++}$, assuming that $\rho_A < 0$, or equivalently, $\ker(A)_{++} \neq \emptyset$.

A remark on the condition numbers is in order. We assume w.l.o.g. that the input matrix A has the property that $\|a_j\| \le 1$ for all $j \in [n]$, and thus $|\rho_A| \le |\hat{\rho}_A|$. The volumetric analysis relies on the Goffin measure $|\hat{\rho}_A|$, whereas the termination condition (Lemma 2.6) uses $|\rho_A|$. The running time in Theorem 2.2 is stated in terms of $|\rho_A|$. However, if we run the algorithm with the renormalized matrix A, then $|\rho_A| = |\hat{\rho}_A|$; thus we obtain a running time bound in terms of $|\hat{\rho}_A|$, as claimed in the Introduction.

We use Dunagan-Vempala (DV) updates as the first order method. We maintain a vector $x \in \mathbb{R}^n$, initialized as $x = \vec{e}$; the coordinates x_i never decrease during the algorithm. We maintain $y = Ax$, and a main quantity of interest is the norm $\|y\|^2$. In each iteration of the algorithm, we check whether $\bar{x} = \Pi x$, the projection of x onto $\ker(A)$, is strictly positive. If this happens, then $\bar{x}$ is returned as a feasible solution to (K_{++}).

Every iteration either applies a DV update to x, thus decreasing the norm of $y = Ax$, or performs a rescaling of the matrix A. It follows from (8) that, if in a given iteration there exists $k \in [n]$ such that $\hat{a}_k^\mathsf{T}\hat{y} \le -\varepsilon$, then we obtain a substantial decrease in the norm, namely

$$\|y'\| \le \|y\|\sqrt{1-\varepsilon^2}. \tag{11}$$

The algorithm proceeds with DV updates as long as there exists such a $k \in [n]$. On the other hand, if $\hat{a}_j^\mathsf{T}\hat{y} \ge -\varepsilon$ for all $j \in [n]$, then it follows that $|\hat{\rho}_A| < \varepsilon$, that is, the condition measure of the current matrix A is small. In terms of the polytope P_A defined in (9),

the condition $\hat{a}_j^\mathsf{T}\hat{y} \geq -\varepsilon$ for all $j \in [n]$ implies then P_A is contained in a "narrow strip" of width 2ε, namely $P_A \subseteq \{z \in \mathbb{R}^m : -\varepsilon \leq \hat{y}^\mathsf{T} z \leq \varepsilon\}$. If we replace A with the matrix $A' := (I + \hat{y}\hat{y}^\mathsf{T})A$, then Lemma 2.4 implies that $\mathrm{vol}(P_{A'}) \geq 3/2\mathrm{vol}(P_A)$. Geometrically, A' is obtained by applying to the columns of A the linear transformation that "stretches" them by a factor of two in the direction of $\hat{y}$ (see Figure 1).

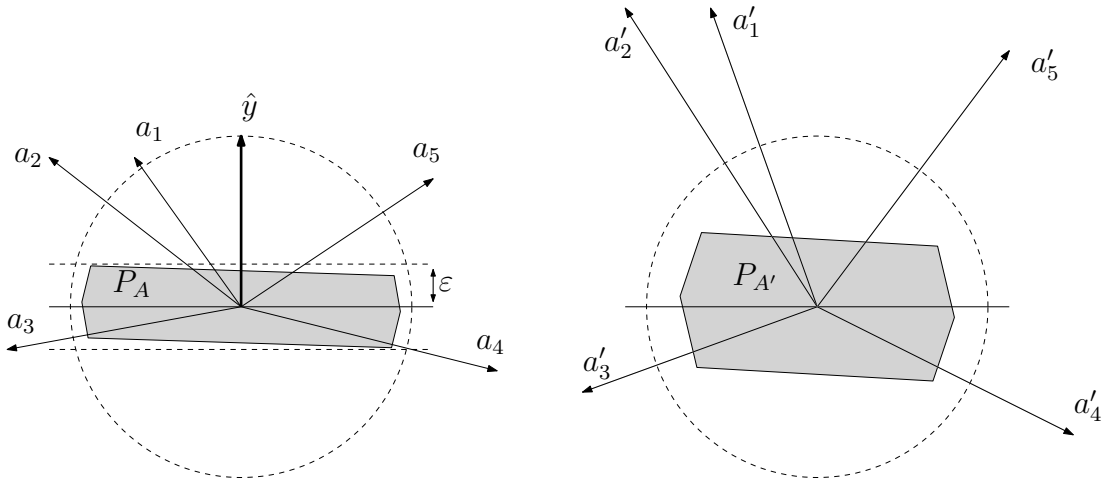

Figure 1: Effect of rescaling. The dashed circle represent the set of points of norm 1. The shaded areas are P_A and $P_{A'}$.

Thus, at every iteration we either have a substantial decrease in the length of the current y, or we have a constant factor increase in the volume of P_A.

The volume of P_A is bounded by the volume of the unit ball in $\mathbb{R}^m$, and initially contains a ball of radius $|\hat{\rho}_A|$ around the origin. Consequently, the number of rescalings cannot exceed $m \log_{3/2} |\hat{\rho}_A|^{-1} \leq m \log_{3/2} |\rho_A|^{-1}$.

The norm $\|y\|$ changes as follows. In every iteration where the DV update is applied, the norm of $\|y\|$ decreases by a factor $\sqrt{1-\varepsilon^2}$ according to (11). At every rescaling, the norm of $\|y\|$ increases by a factor 2. Lemma 2.6 shows that once $\|y\| < |\rho_A|$ for the initial value of $|\rho_A|$, then the algorithm terminates with $\bar{x} = \Pi x > 0$.

Theorem 2.2. *For any input matrix $A \in \mathbb{R}^{m\times n}$ such $\rho_A < 0$ and $\|a_j\| \leq 1$ for all $j \in [n]$, Algorithm 1 finds a feasible solution of (K_{++}) in $O(m^2 \log n + m^3 \log |\rho_A|^{-1})$ DV updates. The number of arithmetic operations is $O(m^2 n \log n + (m^3 n + mn^2) \log |\rho_A|^{-1})$.*

We prove the statement of Theorem 2.2 in the next subsection (Section 2.1.1). Using Lemma 1.4, we obtain a running time bound in terms of bit complexity.

Corollary 2.3. *Let A be an $m\times n$ matrix with integer entries and encoding size L. If $\rho_A < 0$, then Algorithm 1 applied to $\hat{A}$ finds a feasible solution of (K_{++}) in $O\left((m^3 n + mn^2)L\right)$ arithmetic operations.*

Coordinate Descent with Finite Convergence Before proceeding to the analysis, let us consider a modification of Algorithm 1 without any rescaling. This yields a new "pure" coordinate descent method for (K_{++}) in case $\rho_A < 0$ with finite convergence, which was not previously noted in the literature. Again, we assume $\|a_j\| \leq 1$ for all $j \in [n]$. It follows by (11) that the norm $\|y\|$ decreases by at least a factor $\sqrt{1-\rho_A^2}$ in every DV update. Initially, $\|y\| \leq n$, and, as shown in Lemma 2.6, once $\|y\| < |\rho_A|$, the algorithm terminates with a solution $\Pi_A^K x > 0$. Hence the total number of DV steps is bounded by $O(\log(n/|\rho_A|)/\rho_A^2)$.

2.1.1 Analysis

The crucial part of the analysis is to bound the volume increase of P_A at every rescaling iteration.

Lemma 2.4. *Let $A \in \mathbb{R}^{m\times n}$ and let $r = \mathrm{rk}(A)$. For some $0 < \varepsilon < 1/(11r)$, let $v \in \mathbb{R}^m$, $\|v\| = 1$, such that $\hat{a}_j^\mathsf{T} v \geq -\varepsilon\ \forall j \in [n]$. Let $T = (I + vv^\mathsf{T})$, and let $A' = TA$. Then*

i) $TP_A \subseteq (1+3\varepsilon)P_{A'}$,

ii) If $v \in \mathrm{im}(A)$, then $\mathrm{vol}_r(P_{A'}) \geq \frac{3}{2}\mathrm{vol}_r(P_A)$.

The following easy technical claim will be needed. The proof is deferred to the Appendix.

Lemma 2.5. *Let $X \in \mathbb{R}$ be a random variable supported on the interval $[-\varepsilon, \eta]$, where $0 \leq \varepsilon \leq \eta$, satisfying $\mathbb{E}[X] = \mu$. Then for $c \geq 0$, we have that*

$$\mathbb{E}[\sqrt{1+cX^2}] \leq \sqrt{1+c\eta(\varepsilon+|\mu|)}$$

Proof of Lemma 2.4.

i) The statement is trivial if $P_A = \emptyset$, thus we assume $P_A \neq \emptyset$. Consider an arbitrary point $z \in P_A$. By symmetry, it suffices to show $Tz \in (1+3\varepsilon)\mathrm{conv}(\hat{A}')$. By definition, there exists $\lambda \in \mathbb{R}^n_+$ such that $\sum_{j=1}^n \lambda_j = 1$ and $z = \sum_{j=1}^n \lambda_j \hat{a}_j$. Note that

$$Tz = \sum_{j=1}^n \lambda_j T\hat{a}_j = \sum_{j=1}^n (\lambda_j \|T\hat{a}_j\|)\hat{a}'_j = \sum_{j=1}^n \lambda_j \sqrt{1+3(v^\mathsf{T}\hat{a}_j)^2}\,\hat{a}'_j.$$

Since $P_{A'} \neq \emptyset$, it follows that $0 \in \mathrm{conv}(\hat{A}')$, thus it suffices to show that $\sum_{j=1}^n \lambda_j \sqrt{1+3(v^\mathsf{T}\hat{a}_j)^2} \leq 1+3\varepsilon$.

The above is of the form $\mathbb{E}[\sqrt{1+3X^2}]$ where X is a random variable supported on $[-\varepsilon, 1]$ and $|\mathbb{E}[X]| = |\sum_{j=1}^n \lambda_j v^\mathsf{T}\hat{a}_j| = |v^\mathsf{T} z|$. Note that $|v^\mathsf{T} z| \leq \varepsilon$ because both z and $-z$ are in P_A. Hence, by Lemma 2.5,

$$\sum_{j=1}^n \lambda_j \sqrt{1+3(v^\mathsf{T}\hat{a}_j)^2} \leq \sqrt{1+3(2\varepsilon)} \leq 1+3\varepsilon.$$

ii) Note that $\mathrm{vol}_r(TP_A) = 2\mathrm{vol}_r(P_A)$ as $\det(T) = 2$. Thus we obtain $\mathrm{vol}_r(P'_A) \geq 2\mathrm{vol}_r(P_A)/(1+3\varepsilon)^r \geq \frac{3}{2}\mathrm{vol}_r(P_A)$, since $(1+3\varepsilon)^r \leq (1+3/(11m))^r \leq e^{3/11} \leq \frac{4}{3}$. ☐

Lemma 2.6. *Let $A \in \mathbb{R}^{m\times n}$. Given $x \in \mathbb{R}^n$ such that $x \geq \vec{e}$, if $\|Ax\| < |\rho_{(A,-A)}|$ then $\Pi_A^K x > 0$. If $\rho_A < 0$ then $|\rho_{(A,-A)}| \geq |\rho_A|$.*

Proof. Let $\Pi := \Pi_A^K$ and define $\delta \stackrel{\mathrm{def}}{=} \min_{j\in[n]} \|(AA^\mathsf{T})^+ a_j\|^{-1}$. Observe that, if $\|Ax\| < \delta$, then $\Pi_A^K x > 0$. Indeed, for $i \in [n]$, $(\Pi x)_i = x_i - a_i^\mathsf{T}(AA^\mathsf{T})^+ y \geq 1 - \|(AA^\mathsf{T})^+ a_i\|\|Ax\| > 1 - \delta^{-1}\delta = 0$, as required. Thus it suffices to show that $\delta \geq |\rho_{(A,-A)}|$. Since $\mathrm{conv}(A,-A)$ is symmetric around 0, we get that

$$|\rho_{(A,-A)}| = -\rho_{(A,-A)} = \min_{y\in\mathrm{im}(A)\setminus\{0\}} \max_{i\in[n]} |a_i^\mathsf{T}\hat{y}|.$$

Let $k := \arg\max_{j\in[n]} \|(AA^\top)^+ a_j\|$, and plug in $y = (AA^\mathsf{T})^+ a_k$ in the above formula. For this choice, observe that $\|y\| = 1/\delta$ by the choice of k, and therefore

$$|\rho_{(A,-A)}| \leq \delta \max_{j\in[n]} |a_j^\mathsf{T} y| = \delta \max_{j\in[n]} |\Pi_{jk}| \leq \delta,$$

where the last inequality follows since the entries of a projection matrix always have absolute value at most 1.

The second claim is an immediate consequence of Lemma 1.2, using that $\mathrm{conv}(A,-A) \supseteq \mathrm{conv}(A)$ and $\mathrm{conv}(A)$ contains 0 in its interior if $\rho_A < 0$. ☐

Lemma 2.7. *Let $v \in \mathbb{R}^m$, $\|v\| = 1$. For any $y, \bar{y} \in \mathbb{R}^m$ such that $y = (I + vv^\mathsf{T})\bar{y}$, we have $\|\bar{y}\| \leq \|y\|$.*

Proof. We have $\|y\|^2 = \|\bar{y} + (v^\mathsf{T}\bar{y})v\|^2 = \|\bar{y}\|^2 + 2(v^\mathsf{T}\bar{y})^2 + (v^\mathsf{T}\bar{y})^2\|v\|^2 \geq \|\bar{y}\|^2$. □

Proof of Theorem 2.2. Denote by $\bar{A}$ the input matrix, and let $\rho := |\rho_{\bar{A}}|$ and $\Pi := \Pi_{\bar{A}}^K$. Denote by A the current matrix during the algorithm, so that, after k rescalings, $A = (I + v_1v_1^\mathsf{T}) \cdots (I + v_kv_k^\mathsf{T})\bar{A}$ for some vectors $v_1, \ldots, v_k \in \mathrm{im}(A)$ with $\|v_i\| = 1$ for $i \in [n]$. Note that $\ker(A) = \ker(\bar{A})$ throughout the algorithm, so $\Pi_A^K = \Pi$ throughout. Let x and $y = Ax$ be the vectors computed in every iteration, and let $\bar{y} := \bar{A}x$ and $\bar{x} = \Pi x$. Lemma 2.6 implies that $\bar{x} > 0$ whenever $\|\bar{y}\| < |\rho_{(\bar{A},-\bar{A})}|$. Lemma 2.7 implies that $\|\bar{y}\| \leq \|y\|$, and $|\rho_{(\bar{A},-\bar{A})}| \geq \rho$ by Lemma 2.6. It follows that the algorithm terminates with the positive solution $\bar{x} > 0$ as soon as $\|y\| < \rho$.

As previously discussed, Lemma 2.4 implies that the number of rescalings cannot exceed $m \log_{3/2} |\rho|^{-1}$. At every rescaling, the norm of $\|y\|$ increases by a factor 2. In every iteration where the DV update is applied, the norm of $\|y\|$ decreases by a factor $\sqrt{1-\varepsilon^2}$ according to (11). Initially, $y = \bar{A}\vec{e}$, therefore $\|y\| \leq n$ since all columns of $\hat{A}$ have norm at most one. Therefore, the number of DV iterations is bounded by κ, where κ is the smallest integer such that

$$n^2(1-\varepsilon^2)^\kappa 4^{m \log_{3/2} |\rho|^{-1}} < \rho^2.$$

Taking the logarithm on both sides, and using the fact that $\log(1-\varepsilon^2) < -\varepsilon^2$, it follows that $\kappa \in O(m^2 \log n + m^3 \log |\rho|^{-1})$.

We can implement every DV update in $O(n)$ time, at the cost of an $O(n^2)$ time preprocessing at every rescaling, as explained next. Let us compute the matrix $F := A^\mathsf{T}A$ and the norms of the columns of A after every rescaling. The matrix F is updated as $F := A^\mathsf{T}(I + \hat{y}\hat{y}^\mathsf{T})^2A = AA^\mathsf{T} + 3(A^\mathsf{T}\hat{y}\hat{y}^\mathsf{T}A) = F + 3zz^\mathsf{T}/\|y\|^2$, which requires time $O(n^2)$. Further, at every DV update, let us update the vectors $z = A^\mathsf{T}y$ and $\bar{x} = \Pi x$.

Using the vector z, we can compute $\arg\min_{j\in[n]} \hat{a}_j^\mathsf{T}\hat{y} = \arg\min_{j\in[n]} z_j/\|a_j\|$ in time $O(n)$ at any DV update. We also need to recompute y, z, and $\bar{x}$. Using $F = [f_1, \ldots, f_n]$, these can be obtained as $y := y - (\hat{a}_k^\mathsf{T}y)\hat{a}_k$, $z := z - f_k(\hat{a}_k^\mathsf{T}y)/\|a_k\|$, and $\bar{x} := \bar{x} - \Pi_k(\hat{a}_k^\mathsf{T}y)/\|a_k\|$. These updates altogether take $O(n)$ time.

Therefore the number of arithmetic operations is $O(n)$ times the number of DV updates plus $O(n^2)$ times the number of rescalings. The overall running time estimate in the theorem statement follows. □

2.2 Maximum support case

The Maximum Support Kernel Algorithm (Algorithm 2) is a modification of Algorithm 1. In this context, it will be convenient to describe the full support algorithm within a different framework: instead of applying the rescaling transformation to the matrix A, we keep changing the scalar product. If $\bar{A}$ is the input matrix, then after t rescalings in Algorithm 1, the current matrix A is of the form $A = M\bar{A}$. Here $M \in \mathbb{R}^{m\times m}$ is the composition of all t linear transformations applied to $\bar{A}$ thus far. This sequence of updates naturally corresponds to the scalar product defined by the matrix $M^\top M \in \mathbb{S}_{++}^m$, although in the description of the algorithm we use $Q = M^\top M/(1+3\varepsilon)^{2t}$ for the sake of convenience in the analysis.

Consider the vector $y = Ax$ in any iteration of Algorithm 1, and let $\bar{y} = \bar{A}x$. Note that $y = M\bar{y}$ and $y/\|y\|_2 = M\bar{y}/((1+3\varepsilon)^t\|\bar{y}\|_Q)$. When computing $\hat{a}_j^\mathsf{T}\hat{y}$, we have $\hat{a}_j^\mathsf{T}\hat{y} = \left\langle \frac{a_j}{\|a_j\|_Q}, \frac{\bar{y}}{\|\bar{y}\|_Q} \right\rangle_Q$. Furthermore, when rescaling, Algorithm 1 replaces A by $(I_m + \hat{y}\hat{y}^\top)A$, therefore the corresponding update of the scalar product consists of replacing M by $\left(I_m + M\bar{y}(M\bar{y})^\mathsf{T}/\|\bar{y}\|_Q^2\right)M$ and recomputing Q. In terms of Q, the update can

be written as

$$\frac{1}{(1+3\varepsilon)^{2(t+1)}} M^{\mathsf{T}} \left(I_m + \frac{M\bar{y}(M\bar{y})^{\mathsf{T}}}{\|\bar{y}\|_Q^2} \right) \left(I_m + \frac{M\bar{y}(M\bar{y})^{\mathsf{T}}}{\|\bar{y}\|_Q^2} \right) M = \frac{1}{(1+3\varepsilon)^2} \left(Q + \frac{3Q\bar{y}\bar{y}^{\mathsf{T}}Q}{\|\bar{y}\|_Q^2} \right).$$

We remark that the normalization of Q by $(1+3\varepsilon)^{2t}$ can be omitted in the implementation.

Let us now highlight the key ideas of the maximum support setting. Assume the input matrix A has $\|a_i\| \leq 1$ for all $i \in [n]$. We modify A by removing columns that we were able to infer to be outside S^*. We remark that, by Lemma 2.1, P_A does not change when we remove columns from A that are not in S^*. From Lemma 2.4(i), we can easily derive Lemma 2.10, stating that the invariant

$$P_A \subseteq E_{\mathrm{im}(A)}(Q) = \{z \in \mathrm{im}(A) : \|z\|_Q \leq 1\}$$

is maintained at every rescaling (the normalization by $(1+3\varepsilon)^{2t}$ in the definition of Q is needed for this claim). We then show that $\|a_j\|_Q \leq 1/\theta_A$ must hold throughout the algorithm for every $i \in S^*$, where θ_A is the symmetry measure defined in (4). Hence if $\|a_j\|_Q > 1/\theta_A$ at any point of the algorithm, we can infer $i \notin S^*$, and remove the column a_i from A. A volume argument, involving the condition number ρ_A^* defined in (5), shows that within $O(m \log(\theta_A^{-1} {\rho_A^*}^{-1}))$ rescalings, either the algorithm must terminate, or the dimension of the set $\{a_j : \|a_j\|_Q \leq \theta_A^{-1}, j \in [n]\}$ must decrease by one. An amortized argument will show that $O(m \log(\theta_A^{-1} {\rho_A^*}^{-1}))$ will in fact bound the total number of rescalings throughout the algorithm.

In Algorithm 2, we maintain a set $S \subseteq [n]$ with the property that $S^* \subseteq S$, where $S := [n]$ at the beginning. At every iteration, the current matrix A is the matrix obtained from the input matrix by removing all columns outside of S. We also maintain a set $T \subseteq S$ of indices that we have determined not to belong to S^* (that is, $T \cap S^* = \emptyset$ throughout the algorithm). Whenever we conclude that $i \notin S^*$ for an index i based on Lemma 2.10(ii), we add i to T. Columns indexed by T are removed from the current matrix A whenever doing so decreases the rank of the matrix. The algorithm either terminates with a solution $\bar{x}$ to the system $A\bar{x} = 0$, $\bar{x} \geq 0$ with $\mathrm{supp}(\bar{x}) = S$, in which case we may conclude $S = S^*$, or $S = \emptyset$ is reached, in which case we may conclude that $\bar{x} = 0$ is a maximum support solution.

Theorem 2.8. *Let $A \in \mathbb{R}^{m \times n}$ such that $\|a_j\| \leq 1$ for all $j \in [n]$, and let $\bar{\theta}$ be a lower-bound on θ_A. Algorithm 2 finds a solution of $Ax = 0$, $x \geq 0$ of maximum support in $O\left((m^3 n + mn^2) \log(\bar{\theta}^{-1} {\rho_A^*}^{-1})\right)$ arithmetic operations.*

By Lemma 1.4, we can choose $\bar{\theta} \in 1/2^{O(L)}$, and we have $\rho_A^* \in 1/2^{O(L)}$. This implies the following corollary, when applying the algorithm for $\hat{A}$.

Corollary 2.9. *Let A be an $m \times n$ matrix with integer entries and encoding size L. Algorithm 2 finds a solution of $Ax = 0$, $x \geq 0$ of maximum support in $O\left((m^3 n + mn^2)L\right)$ arithmetic operations.*

2.2.1 Analysis

Lemma 2.10. *Throughout the algorithm, the following invariants are maintained.*

(i) $P_A \subseteq E_{\mathrm{im}(A)}(Q) = \{z \in \mathrm{im}(A) : \|z\|_Q \leq 1\}$,

(ii) $\|a_i\|_Q \leq 1/\theta_A$ *for every* $i \in S^*$,

(iii) $S^* \subseteq S$.

Algorithm 2 Maximum Support Kernel Algorithm

Input: A matrix $A \in \mathbb{R}^{m\times n}$ such that $\|a_j\| \le 1$ for all $j \in [n]$, and a lower-bound $\bar{\theta}$ on θ_A.

Output: A maximum support solution to the system (K).

Compute $\Pi := \Pi_A^K$.
Set $x_j := 1$ for all $j \in [n]$, and $y := Ax$.
Set $S := [n]$, $T := \emptyset$, $Q := I_m$.
while $(S \neq \emptyset)$ and $(\Pi x \not> 0)$ **do**
    if $\langle a_i, y\rangle_Q \ge 0$ for all $i \in S$ **then**
        $T := T \cup \{i : \langle a_i, y\rangle_Q > 0\}$, REMOVE$(T)$;
    else
        Let $k := \arg\min_{i\in S} \langle a_i, y\rangle_Q / \|a_i\|_Q$;
        if $\langle a_k, y\rangle_Q < -\varepsilon \|a_k\|_Q \|y\|_Q$ **then**
            update $x := x - \dfrac{\langle a_k, y\rangle_Q}{\|a_k\|_Q^2}\vec{e}_k$; $\quad y := y - \dfrac{\langle a_k, y\rangle_Q}{\|a_k\|_Q^2} a_k$;
        else
            rescale $Q := \dfrac{1}{(1+3\varepsilon)^2}\left(Q + \dfrac{3Qyy^{\mathsf{T}}Q}{\|y\|_Q^2}\right)$;
            if $\exists k \in S \setminus T$ such that $\|\hat{a}_k\|_Q > \bar{\theta}^{-1}$ **then**
                $T := T \cup \{k \in S \setminus T : \|\hat{a}_k\|_Q > \bar{\theta}^{-1}\}$,
                if $\mathrm{rk}(A_{S\setminus T}) < \mathrm{rk}(A)$ **then** REMOVE(T);
if $\Pi x > 0$ **then**
    Define $\bar{x}_i \in \mathbb{R}^n$ by $\bar{x}_i := (\Pi x)_i$ if $i \in S$, $\bar{x}_i := 0$ if $i \notin S$. **return** $\bar{x}$.
if $S = \emptyset$ **then return** $\bar{x} = 0$.

Algorithm 3 Column deletion

procedure REMOVE(T)
    Delete all a_i with $i \in T$ from A. $S := S \setminus T$; $T := \emptyset$.
    Reset $x_j := 1$ for all $j \in S$; $\quad y := Ax$.
    Recompute $\Pi := \Pi_A^K$;

Proof. The three conditions are all trivially satisfied at initialization, since $Q = I_m$, so $P_A \subseteq E(Q) = \mathbb{B}^m$, $\|a_i\|_Q = \|a_i\|_2 = 1 \le 1/\theta_A$ for all $i \in S^*$, and $S = [n]$. Also, as long as invariant (i) is maintained, invariants (ii) and (iii) also hold. Indeed, by definition (4) of θ_A, it follows that $\theta_A a_j \in \mathrm{conv}(A) \cap \mathrm{conv}(-A) \subseteq P_A \subset E(Q)$ for all $j \in S^*$, where the first containment follows from the fact that $\|a_j\| \le 1$ for all $j \in [n]$. Consequently, $\|a_j\|_Q \le 1/\theta_A$ holds for every $i \in S^*$. Finally, since a column a_j is removed only if $\|a_j\|_Q > 1/\theta_A$, we also have $S^* \subseteq S$.

Thus, we only need to show that, if at a given iteration we have $S^* \subseteq S$, then the algorithm will maintain the condition $P_A \subseteq E_{\mathrm{im}(A)}(Q)$ at the subsequent iteration. By Lemma 2.1, $P_A = P_{A_{S^*}}$, so P_A does not change throughout the algorithm since we maintain $S^* \subseteq S$. It suffices to show that $P_A \subseteq E(Q)$, since $P_A \subseteq \mathrm{im}(A_{S^*}) \subseteq \mathrm{im}(A)$.

After t rescalings, let M and Q be the corresponding matrices, so that $Q = (1+\varepsilon)^{-2t} M^{\mathsf{T}} M$. Let $A' = MA$. By Lemma 2.4 applied t times, $MP_A \subseteq (1+3\varepsilon)^t P_{A'}$. It follows that $P_A \subseteq (1+3\varepsilon)^t M^{-1}\mathbb{B}^m = E(Q)$. ∎

Lemma 2.11. *The relative volume of the ellipsoid $E_{\mathrm{im}(A)}(Q)$ decreases at least by a factor $3/2$ at every rescaling.*

Proof. At rescaling, we updated the matrix Q as $Q' = (1+3\varepsilon)^{-2}(Q + 3Qyy^{\mathsf{T}}Q/\|y\|_Q^2))$, where $y \in \mathrm{im}(A)$ for the current matrix A. Let $H = \mathrm{im}(A)$. We need to prove that $\det_H(Q') \geq (9/4)\det_H(Q)$. Since $(1+3\varepsilon)^m \leq 4/3$, it suffices to show that $\det_H(Q + 3Qyy^{\mathsf{T}}Q/\|y\|_Q^2) = 4$. Let W be an $m \times \mathrm{rk}(A)$ matrix whose columns for an orthonormal basis of $\mathrm{im}(A)$. Since $y \in \mathrm{im}(A)$, we have $y = Wz$ for some $z \in \mathbb{R}^{\mathrm{rk}(A)}$. Then

$$\begin{aligned}\det_H(Q + 3Qyy^{\mathsf{T}}Q/\|y\|_Q^2) &= \det(W^{\mathsf{T}}QW + 3W^{\mathsf{T}}QWzz^{\mathsf{T}}W^{\mathsf{T}}QW/\|y\|_Q^2)\\ &= \det(W^{\mathsf{T}}QW)(1 + 3z^{\mathsf{T}}W^{\mathsf{T}}QWz/\|y\|_Q^2)\\ &= \det(W^{\mathsf{T}}QW)(1 + 3y^{\mathsf{T}}Qy/\|y\|_Q^2) = 4\end{aligned}$$

where the second equality follows from the Sherman-Morrison formula (that is, for every non-singular $B \in \mathbb{R}^{m\times m}$ and $u, v \in \mathbb{R}^m$, $\det(B + uv^{\mathsf{T}}) = \det(B)(1 + v^{\mathsf{T}}B^{-1}u)$). □

We now formulate the two main technical lemmas, and derive the proof of Theorem 2.8 using them. The proof of these two lemmas is deferred to Section 2.2.2. In the following, $\bar{\theta}$ denotes the known lower bound on the initial θ_A, and ρ^* refers to the initial ρ_A^*.

Lemma 2.12. *At every call of the procedure* REMOVE*, the rank of A decreases. The algorithm can be implemented so that the total number of operations required for all calls to the procedure throughout the execution of Algorithm 2 is bounded by $O(mn^2)$.*

Lemma 2.13. *Let A' be the matrix obtained from the current matrix A after a call of the procedure* REMOVE*. Let $r' := \mathrm{rk}(A') < r$, $E := E_{\mathrm{im}(A)}(Q)$, and $E' := E_{\mathrm{im}(A')}(Q)$. Then*

$$\frac{\mathrm{vol}_{r'}(E')}{\mathrm{vol}_r(E)} \leq \frac{\nu_{r'}}{\nu_r}\left(\frac{2}{\bar{\theta}\rho^*}\right)^{r-r'}.$$

Proof of Theorem 2.8. From Lemma 2.10, whenever Algorithm 2 returns a solution, it is a maximum support solution. We need to argue that the algorithm terminates in the prescribed number of iterations. Let $\bar{A}$ denote the input matrix, and A be the current matrix at any given iteration. Recall that $A = \bar{A}_S$ for some $S \supseteq S^*$, and that we have $P_A = P_{\bar{A}}$ by Lemma 2.1. We let r denote the rank of the current matrix A. Also recall that the input requires $\|\bar{a}_j\| \leq 1$ for all $j \in [n]$.

We define a *round* of the algorithm the iterations that take place between two iterations in which we remove some columns from the matrix A. During each round, the algorithm maintains the set $T = \{j \in S : \|a_j\|_Q > \bar{\theta}^{-1}\}$ of columns eligible for removal. Columns are removed if $\mathrm{rk}(A_{S\setminus T}) < \mathrm{rk}(A)$. In particular, the number of rounds is at most $\mathrm{rk}(A) \leq m$. We want to bound the total number of rescalings performed by the algorithm.

Claim 2.14. *Thus the total number of rescalings throughout the algorithm is $O(m\log(\bar{\theta}^{-1}\rho^{*-1}))$.*

Proof. We will show next that at every rescaling within the same round, except for the last rescaling of the round, the invariant

$$\mathrm{vol}_r(E) \geq \nu_r(\bar{\theta}\rho^*)^r \tag{12}$$

is maintained. Indeed, by Lemma 2.10(i), $P_A \subseteq E := E_{\mathrm{im}(A)}(Q)$ throughout. Since $\|a_j\|_Q \leq \bar{\theta}^{-1}$ for all $j \in S \setminus T$, it follows that $E \supseteq \bar{\theta}\mathrm{conv}(A_{S\setminus T}, -A_{S\setminus T})$. Since at every rescaling except for the last one of the round we have $\mathrm{rk}(A_{S\setminus T}) = r$, it follows by Lemma 1.2 that $\mathrm{conv}(A_{S\setminus T}, -A_{S\setminus T})$ contains an r-dimensional ball of radius $\rho_{(A_{S\setminus T}, -A_{S\setminus T})} \geq \rho^*$. This implies (12).

At the first iteration, $Q = I_m$ and $E = \mathbb{B}_m \cap \mathrm{im}(\bar{A})$, therefore initially $\mathrm{vol}_{\bar{r}}(E) \leq \nu_{\bar{r}}$, where $\bar{r} = \mathrm{rk}(\bar{A})$. By Lemma 2.11, at every rescaling in which we do not remove any

column $\mathrm{vol}_r(E)$ decreases by at least 2/3; Lemma 2.13 bounds the increase in $\mathrm{vol}_r(E)$ at column removals. Combined with the lower bound (12), we obtain that the total number of rescalings is at most m plus the smallest number K satisfying

$$(2/3)^K \left(\frac{2}{\bar{\theta}\rho^*}\right)^m < (\bar{\theta}\rho^*)^m.$$

The claimed bound on K follows. □

By Lemma 2.6, the algorithm is guaranteed to terminate when $\|y\|_2 < |\rho_{(A,-A)}| \leq \rho^*$. By Lemma 2.7, after t rescalings $\|y\|_2 \leq \|y\|_Q(1+3\varepsilon)^t$. Hence the algorithm is guaranteed to terminate if $\|y\|_Q \leq \rho^*/(1+3\varepsilon)^K$.

At the beginning of each round, we re-initialize x so that $x_j = 1$ for all $j \in S$. In particular, $y = Ax$ satisfies $\|y\|_Q \leq |S| \max_{j\in[S]} \|a_j\|_Q \leq n\bar{\theta}^{-1}$, where the last inequality follows from $\|a_j\|_Q \leq \bar{\theta}^{-1}$ for all $j \in S$. At every rescaling within the same round, $\|y\|_Q$ increases by $2/(1+3\varepsilon)$, and in every DV update, it decreases by at least a factor $\sqrt{1-\varepsilon^2}$. Let T be the number of rounds, and $K_1, \ldots, K_T$ be the number of rescaling within rounds $1, \ldots, T$.

It follows that, at the ith round, the number of DV updates is at most the smallest number κ_i such that

$$n\bar{\theta}^{-1}(1-\varepsilon^2)^{\kappa_i/2}2^{K_i} < \rho^*/(1+3\varepsilon)^K.$$

Taking the logarithm on both sides and recalling that $\log(1-\varepsilon^2) < -\varepsilon^2$ and $\log(1+3\varepsilon) \geq 3\varepsilon$, it follows from our choice of ε that $\kappa_i \in O(m^2)K_i + O(m)K$. Since $K = K_1 + \cdots + K_T$ and $T \leq m$, this implies that the total number of DV updates is $O(m^2)K$. Using Claim 2.14, the total number of DV updates is $O(m^3 \log(\bar{\theta}\rho^*)^{-1})$. As explained in the proof of Theorem 2.2, we can perform each DV update in $O(n)$ arithmetic operations, provided that at each rescaling we recompute $F = A^{\mathsf{T}}QA$ at a cost of $O(n^2)$ arithmetic operations. This implies the stated running time. □

2.2.2 Proofs of technical lemmas

Proof of Lemma 2.12. Algorithm 2 removes columns from the matrix A in two cases (in step 5 and step 13) of the algorithm. In step 13 we remove the columns in T only if this decreases the rank of the matrix A. We show that after the column-removal in step 5, the rank of A decreases. This will imply the first part of the statement. Indeed, in step 5 we determined x with $\|x\|_1 = 1$ such that $0 \neq y = Ax$ and $A^{\mathsf{T}}Qy \geq 0$. Since $0 < \|y\|_Q^2 = x^{\mathsf{T}}A^{\mathsf{T}}Qy$, it follows that $A^{\top}Qy \neq 0$. This shows that there exists $k \in S$ such that $\langle a_k, y\rangle_Q > 0$, so in particular $S' := \{j \in S : \langle a_j, y\rangle_Q = 0\} \subset S$, which implies that $\mathrm{rk}(A_{S'}) < \mathrm{rk}(A)$.

At line 13 we need to check if $\mathrm{rk}(A_{S\setminus T}) < \mathrm{rk}(A)$. We next observe how this check can be implemented so that the total number of arithmetic operations carried out for this purpose throughout the entire execution of the algorithm is $O(mn^2)$. Throughout the algorithm, while $\mathrm{rk}(A_{S\setminus T}) = \mathrm{rk}(A)$ we maintain a set $B \subseteq S \setminus T$ of indices of $\mathrm{rk}(A)$ linearly independent columns of A, and a $\mathrm{rk}(A) \times |S|$ matrix F that is obtained from A by Gaussian elimination so that the columns of F indexed by B form an identity matrix. At the beginning of the algorithm (in which case $S := [n]$ and $T = \emptyset$) this will require $O(m^2n)$ arithmetic operations. Suppose now that, in a given iteration, for $k \in S \setminus T$ we want to check if removing from A the columns in $T' := T \cup \{k\}$ decreases the rank of the current matrix A. If $k \notin B$, then the rank does not decrease, and we maintain the property that $B \cap T' = \emptyset$. If $k \in B$, then let h be the index corresponding to the entry of value 1 in the unit column indexed by k in F and check if there exists an index $j \in S \setminus (B \cup T)$ such that the (h, j) entry of F has nonzero value. If such index j does not exist, then $\mathrm{rk}(A_{S\setminus T'}) < \mathrm{rk}(A)$, so we set $B := B \setminus \{k\}$ and remove from F the columns indexed by T' and the zero row thus created. Otherwise, update $B := B \setminus \{k\} \cup \{j\}$ and

recompute F by performing a pivot on entry (h, j). Performing the pivot requires $O(mn)$ time. Since we remove at most n columns throughout the algorithm, the total running time required is $O(mn^2)$. □

The next three lemmas will also be needed in the proof of Lemma 2.13, as well as in Sections 3. Given a convex set $X \subset \mathbb{R}^d$ and a vector $a \in \mathbb{R}^d$, we define the *width of X along a* as

$$\mathrm{width}_X(a) \stackrel{\text{def}}{=} \max\{a^{\mathsf{T}} z : z \in X\}. \tag{13}$$

Lemma 2.15. *Let $E := E(D^{-1}) = \{x \in \mathbb{R}^d : x^\top D^{-1} x \leq 1\}$, where D is a symmetric positive definite matrix. Given a vector $a \in \mathbb{R}^d$, $\mathrm{width}_E(a) = \|a\|_D$.*

Proof. For every $z \in E$, $a^{\mathsf{T}} z = a^{\mathsf{T}} D^{1/2} D^{-1/2} z \leq \|a\|_D \|z\|_{D^{-1}} \leq \|a\|_D$ where the first inequality follows from the Cauchy-Schwarz inequality and the second from $z \in E$. On the other hand, if we define $z = Da/\|a\|_D$, it follows that $z \in E$ and $a^{\mathsf{T}} z = \|a\|_D$. □

Lemma 2.16. *Let $E \subset \mathbb{R}^d$ be an ellipsoid. Given $a \in \mathbb{R}^d$, $a \neq 0$, let $H = \{x : a^\top x = 0\}$. Then*

$$\mathrm{vol}_{d-1}(E \cap H) = \frac{\mathrm{vol}_d(E)}{\mathrm{width}_E(\hat{a})} \cdot \frac{\nu_{d-1}}{\nu_d}.$$

The lemma follows easily from the next bound on determinants, which we will also need in Section 3.

Lemma 2.17. *Consider a matrix $R \in \mathbb{S}^d_{++}$. For a vector $a \in \mathbb{R}^d$, $a \neq 0$, let $H = \{x : a^\top x = 0\}$. Then*

$$\det_H(R) = \det(R) \|\hat{a}\|^2_{R^{-1}}.$$

Proof. Without loss of generality we can assume that $\|a\|_2 = 1$, that is, $\hat{a} = a$. Let $W \in \mathbb{R}^{d \times (d-1)}$ be a matrix whose columns form an orthonormal basis of H. Since $(W|a)$ is an orthonormal basis of $\mathbb{R}^d$, we have

$$\begin{aligned} \det(R) &= \det\left(\begin{pmatrix} W^\top \\ a^\top \end{pmatrix} R(W|a)\right) = \det\begin{pmatrix} W^\top RW & W^\top Ra \\ a^\top RW & \|a\|^2_R \end{pmatrix} \\ &= \det(W^\top RW)(\|a\|^2_R - a^\top RW(W^\top RW)^{-1} W^\top Ra), \end{aligned}$$

where the last equality follows from the determinant identity for the Shur's complement. We now observe that $\|a\|^2_R - a^\top RW(W^\top RW)^{-1} W^\top Ra = \|q\|^2$, where q is the orthogonal projection of the vector $R^{\frac{1}{2}} a$ onto the orthogonal complement of the hyperplane $R^{1/2} H$. The orthogonal complement of this hyperplane is the line generated by $p = R^{-1/2} a$. Thus,

$$\|q\|^2 = a^\top R^{\frac{1}{2}} R^{-\frac{1}{2}} a (a^\top R^{-1} a)^{-1} a^\top R^{-\frac{1}{2}} R^{\frac{1}{2}} a = \|a\|^2_{R^{-1}}.$$

Therefore $\det(W^\top RW) = \det(R)\|a\|^2_{R^{-1}}$, as required. □

Proof of Lemma 2.16. The volume of E can be written as $\mathrm{vol}_d(E) = \nu_d / \sqrt{\det(R)}$, and using (2), we get $\mathrm{vol}_{d-1}(E \cap H) = \nu_{d-1} / \sqrt{\det_H(R)}$. The statement follows from Lemmas 2.15 and 2.17. □

Proof of Lemma 2.13. It suffices to prove the claim for the case $r' = r - 1$, therefore we assume that $\mathrm{im}(A')$ is a hyperplane of $\mathrm{im}(A)$. Let $a \in \mathrm{im}(A)$ be a vector orthogonal to $\mathrm{im}(A')$ such that $\|a\|_2 = 1$. By Lemma 2.16,

$$\mathrm{vol}_{r'}(E') = \frac{\mathrm{vol}_r(E)}{\mathrm{width}_E(a)} \cdot \frac{\nu_{r-1}}{\nu_r}.$$

Let $S' := \{j \in S : \|a_j\|_Q \leq 2\bar{\theta}^{-1}\}$. Since at every rescaling the Q-norm of the columns of $\hat{A}$ increases by at most a factor 2, the columns in $T \setminus S'$ had Q-norm greater than $\bar{\theta}^{-1}$

already at the previous rescaling. Since we did not remove the columns in $T \setminus S'$ at the previous rescaling, it follows that $\mathrm{rk}(A_{S'}) = \mathrm{rk}(A)$, so $\mathrm{im}(A_{S'}) = \mathrm{im}(A)$. In particular $a \in \mathrm{im}(A_{S'})$, therefore we have

$$\begin{aligned}
\mathrm{width}_E(a) &= \max\{a^\mathsf{T} z : \|z\|_Q \le 1,\, z \in \mathrm{im}(A)\} \\
&\ge \max_{k \in S'} \frac{|a^\mathsf{T} a_k|}{\|a_k\|_Q} \ge \frac{\bar\theta}{2} \min_{y \in \mathrm{im}(A_{S'})} \max_{k \in S'} |\hat y^\mathsf{T} a_k| \\
&= \frac{\bar\theta}{2} |\rho_{(A_{S'}, -A_{S'})}| \ge \frac{\bar\theta \rho_A^*}{2},
\end{aligned}$$

where the last inequality is by definition of ρ_A^*. □

3 The Image Algorithm

We now present our algorithms for the image space. Again, we start by describing the full support version first (Section 3.2), followed by the maximum support version (Section 3.3). In contrast to the kernel algorithms, these two versions will use the same framework of changing the scalar product, and even the same condition measure $\hat\omega_A$. The maximum support version is obtained by a direct extension of the full support version.

The full support algorithm can be seen as an improved version of previous algorithms by Betke [4]. While we use a similar volumetric potential, a more sophisticated rescaling allows for a factor n improvement in the overall number of iterations. In the algorithm we keep modifying the scalar product $\langle ., . \rangle_Q$ via rescalings. For every rescaling, we run the von Neumann algorithm, using the current scalar product $\langle ., . \rangle_Q$, for $O(m^2)$ iterations. This provides a vector $y \in \mathrm{conv}(a_1/\|a_1\|_Q, \ldots, a_n/\|a_n\|_Q)$ with $\|y\|_Q \le \varepsilon$. The coefficients of the convex combination will be used to construct an appropriate rescaling.

Instead of von Neumann's algorithm, one could use other first order methods, such as Perceptron, the DV-updates, or Wolfe's nearest-point algorithm [38]. All these first order methods are interchangeable in this framework, as the only property needed is that they can find a vector $y \in \mathrm{conv}(a_1/\|a_1\|_Q, \ldots, a_n/\|a_n\|_Q)$ with $\|y\| \le 1/\mathrm{poly}(m)$ in time polynomial in m and n. We can use Nesterov's smoothing technique [25] (as illustrated by the Smoothed Perceptron algorithm by Peña and Soheili [34], and the Mirror Prox for Feasibility Problems (MPFP) by Yu et al. [39]).

As mentioned above, our rescaling is an improvement over the one used by Betke [4], who combined his rescaling with a variant of Wolfe's nearest-point algorithm. Peña and Soheili [30] use Betke's rescaling combined with either the Perceptron or the Smoothed Perceptron algorithm. Betke's rescaling requires more first-order iterations, namely, $O(m^2 n)$ calls for Perceptron or Wolfe's nearest-point algorithms; the reason is that his rescaling is based on a single column a_k that has the highest coefficient in the convex combination. To explain the difference, consider an input matrix with unit length columns. For the column a_k with the largest coefficient in the convex combination returned by the first order algorithm, Betke would perform the rescaling $A' := (I_m - \frac{1}{2} a_k a_k^\top) A$. In contrast, our update corresponds to changing the matrix to $A' := (I_m + \sum_{i=1}^n x_i a_i a_i^\mathsf{T})^{-1/2} A$, where $y = Ax$. (Note that in the description of our algorithm we prefer to modify the scalar product instead of changing the matrix.)

We note that our algorithm also has deeper connections to Chubanov's algorithm [10] for the kernel problem; rescaling multiplies A by a diagonal matrix from the right. These connections will be explored in a subsequent paper in the series.

3.1 The von Neumann algorithm

We now state the von Neumann subroutine (Algorithm 4) in the form needed for our algorithm. This is the same as the algorithm described by Dantzig [13], with the scalar

product replaced by $\langle .,. \rangle_Q$ for a matrix $Q \in \mathbb{S}^m_{++}$, and using the normalized columns $a_i/\|a_i\|_Q$.

Algorithm 4 The von Neumann algorithm

Input: A matrix $A \in \mathbb{R}^{m\times n}$, a positive definite matrix $Q \in \mathbb{R}^{m\times m}$ and an $\varepsilon > 0$.
Output: Vectors $x \in \mathbb{R}^n$, $y \in \mathbb{R}^m$ such that $y = \sum_{i=1}^n x_i a_i/\|a_i\|_Q$, $\vec{e}^\top x = 1$, $x \geq 0$, and either $A^\mathsf{T} Qy > 0$ or $\|y\|_Q \leq \varepsilon$.

1: Set $x := \vec{e}_1$, $y := a_1/\|a_1\|_Q$.
2: **while** $\|y\|_Q > \varepsilon$ **do**
3: Let $k := \arg\min_{i\in[n]} \langle a_i, y\rangle_Q / \|a_i\|_Q$;
4: **if** $\langle a_k, y\rangle_Q > 0$ **then return** the solution (x, y) such that $A^\mathsf{T} Qy > 0$. Terminate.
5: **else**
6:

$$\lambda := \frac{\langle y - a_k/\|a_k\|_Q, y\rangle_Q}{\|y - a_k/\|a_k\|_Q\|_Q^2};$$

7: **update** $x := (1-\lambda)x + \lambda\vec{e}_k$; $\quad y := (1-\lambda)y + \lambda\dfrac{a_k}{\|a_k\|_Q}$;

return the vectors (x, y).

The value of λ for the update is choosen so as to minimize the norm $\|(1-\lambda)y + \lambda\frac{a_k}{\|a_k\|_Q}\|_Q$. The following lemma summarizes the properties of the von Neumann algorithm as shown in [13]. Running the algorithm with $\langle .,. \rangle_Q$ is the same as running it for the standard scalar product for the unit vectors $Q^{1/2}a_i/\|Q^{1/2}a_i\|_2$.

Lemma 3.1. *For input $\varepsilon > 0$, the von Neumann algorithm terminates in at most $\lceil 1/\varepsilon^2 \rceil$ updates.*

3.2 Full support case

Algorithm 5 Full Support Image Algorithm

Input: A matrix $A \in \mathbb{R}^{m\times n}$ such that $\mathrm{rk}(A) = m$ and (I_{++}) is feasible.
Output: A feasible solution to (I_{++}).

1: Set $Q := I_m$, $R := I_m$.
2: **while** $A^\mathsf{T} Qy \not> 0$ **do**
3: Call VON NEUMANN(A, Q, ε) to obtain (x, y).
4: **rescale**

$$R := \frac{1}{1+\varepsilon}\left(R + \sum_{i=1}^n \frac{x_i}{\|a_i\|_Q^2} a_i a_i^\top\right); \quad Q := R^{-1}.$$

return The feasible solution $\bar{y} = Qy$ to (I_{++}).

We use the same value $\varepsilon = \frac{1}{11m}$ as in the kernel algorithm. Here as well as in the Maximum Support Image Algorithm, we assume that the matrix A has full row rank, that is, $\mathrm{im}(A) = \mathbb{R}^m$. This is without loss of generality; the general case can be easily

reduced to this case. Let us state the running time bound, which will be proved in Section 3.2.1.

Theorem 3.2. *For any input matrix $A \in \mathbb{R}^{m\times n}$ such that $\mathrm{rk}(A) = m$ and (I_{++}) is feasible, Algorithm 5 finds a feasible solution to (I_{++}) by performing $O\left(m^3 \log \hat{\omega}_A^{-1}\right)$ von Neumann iterations. The number of arithmetic operations is $O\left(m^2 n^2 \log \hat{\omega}_A^{-1}\right)$.*

Using Lemma 1.4, we obtain the running time in terms of the encoding length L.

Corollary 3.3. *Let $A \in \mathbb{Z}^{m\times n}$ be an integer matrix of encoding size L. If $\mathrm{rk}(A) = m$ and (I_{++}) is feasible, then Algorithm 5 finds a feasible solution of (I_{++}) in $O\left(m^2 n^2 L\right)$ arithmetic operations.*

If instead of the von Neumann algorithm we use Nesterov's smoothing technique [25, 34, 39], we obtain the following alternative running-time bound.

Theorem 3.4. *For any input matrix $A \in \mathbb{R}^{m\times n}$ such that $\mathrm{rk}(A) = m$ and (I_{++}) is feasible, Algorithm 5 with Nesterov's smoothing technique finds a feasible solution to (I_{++}) by performing $O\left(m^2\sqrt{\log n} \cdot \log \hat{\omega}_A^{-1}\right)$ iterations. The number of arithmetic operations is $O\left(m^3 n\sqrt{\log n} \cdot \log \hat{\omega}_A^{-1}\right)$. If $A \in \mathbb{Z}^{m\times n}$ is integer with encoding length L, then the running time is $O\left(m^3 n\sqrt{\log n} \cdot L\right)$.*

Oracle model for strict conic feasibility. Let us also remark that the algorithm immediately extends to an oracle model, where the goal is to find a point in the interior of a full-dimensional cone $\Sigma \subseteq \mathbb{R}^m$, given via a separation oracle. That is, for any vector v, the oracle decides whether $v \in \mathrm{int}(\Sigma)$, and if not, it returns a vector z such that $z^\top v \le 0$ but $z^\top y > 0$ for any $z \in \mathrm{int}(\Sigma)$. Hence we can run von Neumann algorithm using this oracle, and use the vectors return by the oracle for rescaling. The algorithm in this case performs $O(m^3 \log(\hat{\rho}^{-1}))$, where $\hat{\rho}$ denotes the radius of the largest ball contained in Σ and centered on the surface of the unit sphere. This model was studied by Peña and Soheili [30], who gave an algorithm that performs $O(m^5 \log(\hat{\rho}^{-1}))$ perceptron (or von Neumann) updates. We will elaborate further on the oracle model in a subsequent paper, in particular, show its application to submodular function minimization.

3.2.1 Analysis

It is easy to see that the matrix R remains positive semidefinite throughout the algorithm, and admits the following decomposition.

Lemma 3.5. *At any stage of the algorithm, we can write the matrix R in the form*

$$R = \alpha I_m + \sum_{i=1}^{n} \gamma_i \hat{a}_i \hat{a}_i^{\mathsf{T}}$$

where $\alpha = 1/(1+\varepsilon)^t$ for the total number of rescalings t performed thus far, and $\gamma_i \ge 0$. The trace is $\mathrm{tr}(R) = \alpha m + \sum_{i=1}^n \gamma_i$.

Recall that we denote by $\Sigma_A = \{y \in \mathbb{R}^m : A^{\mathsf{T}} y \ge 0\}$ the image cone. Let us define the set

$$F_A = \Sigma_A \cap \mathbb{B}^m.$$

The running time will be bounded in terms of the condition number $\hat{\omega}_A$ defined in (6). Using the notation F_A, the definition is $\hat{\omega}_A = \min_{i\in T^*} \mathrm{width}_{F_A}(\hat{a}_i)$. In the full support case, we assume $T^* = [n]$.

The ellipsoid $E(R) = \{z \in \mathbb{R}^m : \|z\|_R^2 \le 1\}$ plays a key role in the analysis, due to the following property:

Lemma 3.6. *Throughout the algorithm, $F_A \subseteq E(R)$ holds.*

The correctness and running time estimation follows by the next three lemmas; Lemma 3.6 is essential in the proof of the first one.

Lemma 3.7. *Throughout the algorithm, $\|\hat{a}_i\|_Q \geq \hat{\omega}_A$ for every $i \in [n]$.*

Lemma 3.8. *The determinant of R increases at least by a factor $16/9$ at every rescaling.*

Lemma 3.9. *At any stage of the algorithm, there exists a column a_k with*

$$\|\hat{a}_k\|_Q \leq \frac{1}{\sqrt{\det(R)^{1/m} - 1}}.$$

We now present the proofs of Theorems 3.2 and 3.4 based on these lemmas.

Proof of Theorem 3.2. After t rescalings, Lemmas 3.8 and 3.9 imply the existence of a column a_k with $\|\hat{a}_k\|_Q \leq 1/(\sqrt{(16/9)^{t/m} - 1})$. Lemma 3.7 yields $1/\hat{\omega}_A \geq \sqrt{(16/9)^{t/m} - 1}$. Consequently, the total number of rescalings during the entire course of the algorithm is providing the bound $O(m \log \hat{\omega}_A^{-1})$.

By Lemma 3.1, the von Neumann algorithm performs $O(m^2)$ iterations between two consecutive rescalings. As in the proof of Theorem 2.2, a von Neumann iteration can be implemented in time $O(n)$, assuming that we compute the matrix $A^\mathsf{T} Q A$ at the beginning and after every rescaling. Provided Q, this can be done in time $O(n^2 m)$. This is more costly than in the kernel case, where we only had to implement rank one updates at every rescaling, since now we need to recompute $A^\mathsf{T} Q A$ from scratch. For every rescaling, this computation dominates the term $O(m^2 n)$ for the von Neumann iterations, as well as the complexity of updating the matrix R and computing the inverse Q. Hence we obtain the overall complexity $O(n^2 m^2 \log \hat{\omega}_A^{-1})$. □

Proof of Theorem 3.4. The smoothed algorithms [25, 34, 39] can either find a $y \in \mathbb{R}^m$ with $A^\mathsf{T} Q y > 0$, or a $y \in \operatorname{conv}(\{a_i/\|a_i\|_Q : i \in [n]\})$ with $\|y\|_Q \leq \varepsilon$ in $O(\sqrt{\log n}/\varepsilon)$ iterations. However, the complexity of each iteration is $O(mn)$ compared to $O(n)$ for the von Neumann algorithm. We obtain the same bound $t = O(m \log \hat{\omega}_A^{-1})$ on the total number of iterations as in the proof above. These together give the claimed bound. Note that we do not have to compute $A^\mathsf{T} Q A$ at every rescaling as it does not yield any computational benefit. □

Lemma 3.6 will be a consequence of the following simple claim.

Claim 3.10. *Let $Q \in \mathbb{S}^m_{++}$, let $R = Q^{-1}$, and let $x \in \mathbb{R}^n_+$ and $y \in \mathbb{R}^m$ such that $y = \sum_{i=1}^n x_i \frac{a_i}{\|a_i\|_Q}$, $\vec{e}^\mathsf{T} x = 1$, $\|y\|_Q \leq \varepsilon$. Then, for every $z \in \Sigma_A$,*

$$\sum_{i=1}^n x_i \left(\frac{a_i^\mathsf{T} z}{\|a_i\|_Q} \right)^2 \leq \varepsilon \|z\|_R^2.$$

Proof. For every $i \in [m]$, we obtain $0 \leq a_i^\mathsf{T} z = a_i^\mathsf{T} Q^{1/2} Q^{-1/2} z \leq \|a_i\|_Q \|z\|_R$ using the Cauchy-Schwartz inequality. Then

$$\sum_{i=1}^n x_i \left(\frac{a_i^\mathsf{T} z}{\|a_i\|_Q} \right)^2 \leq \sum_{i=1}^n x_i \frac{a_i^\mathsf{T} z}{\|a_i\|_Q} \|z\|_R = y^\mathsf{T} z \|z\|_R \leq \|y\|_Q \|z\|_R^2 \leq \varepsilon \|z\|_R^2.$$

The first inequality uses the previous estimate, $x \geq 0$, and $z \in \Sigma_A$; the second inequality uses again the Cauchy-Schwarz inequality. □

Proof of Lemma 3.6. The proof is by induction on the number of rescalings. At initialization, $F_A \subseteq E(I_m) = \mathbb{B}^m$ is trivial. Assume $F_A \subseteq E(R)$, and we rescale R to R'. We show $F_A \subseteq E(R')$. Consider an arbitrary point $z \in F_A$; by the induction hypothesis, $\|z\|_R^2 \leq 1$ because $z \in E(R)$.

For the vector x returned by the von Neumann algorithm, the algorithm sets $R' = \left(R + \sum_{i=1}^n x_i a_i a_i^\mathsf{T}/\|a_i\|_Q^2\right)/(1+\varepsilon)$. Thus Claim 3.10 implies that

$$\|z\|_{R'}^2 = z^\mathsf{T}\left(R + \sum_{i=1}^n \frac{x_i}{\|a_i\|_Q^2} a_i a_i^\mathsf{T}\right) z/(1+\varepsilon) \le \|z\|_R^2$$

Consequently, $z \in E(R')$, completing the proof. □

Proof of Lemma 3.7. Since $F_A \subseteq E(R)$, from Lemmas 2.15 and 3.6 we obtain

$$\|\hat{a}_i\|_Q = \mathrm{width}_{E(R)}(\hat{a}_i) \ge \mathrm{width}_{F_A}(\hat{a}_i) \ge \hat{\omega}_A.$$

The final equality follows since $\hat{\omega}_A = \min_{i\in T^*} \mathrm{width}_{F_A}(\hat{a}_i)$ by definition, and $T^* = [n]$ in case (I_{++}) is feasible. □

Proof of Lemma 3.8. Let R and R' denote the matrix before and after the rescaling. Let $X = \sum_{i=1}^n x_i a_i a_i^\mathsf{T}/\|a_i\|_Q^2$; hence $R' = (R+X)/(1+\varepsilon)$. The ratio of the two determinants is

$$\frac{\det(R')}{\det(R)} = \frac{\det(R+X)}{(1+\varepsilon)^m \det(R)} = \frac{\det\left(I_m + R^{-1/2} X R^{-1/2}\right)}{(1+\varepsilon)^m}$$

Now $R^{-1/2} = Q^{1/2}$, and $Q^{1/2} X Q^{1/2}$ is a positive semidefinite matrix. The determinant can be lower bounded using using Lemma 1.1(iii) and then Lemma 1.1(ii):

$$\frac{\det(R')}{\det(R)} \ge \frac{1 + \mathrm{tr}(Q^{1/2} X Q^{1/2})}{(1+\varepsilon)^m} = \left(1 + \sum_{i=1}^n \frac{x_i}{\|a_i\|_Q^2} \mathrm{tr}(Q^{1/2} a_i a_i^\mathsf{T} Q^{1/2})\right)/(1+\varepsilon)^m.$$

Finally, Lemma 1.1(i) gives $\mathrm{tr}(Q^{1/2} a_i a_i^\mathsf{T} Q^{1/2}) = \mathrm{tr}(a_i^\mathsf{T} Q a_i) = \|a_i\|_Q^2$. Therefore we conclude

$$\frac{\det(R')}{\det(R)} \ge \frac{1 + \sum_{i=1}^n x_i}{(1+\varepsilon)^m} = \frac{2}{(1+\varepsilon)^m}.$$

Using that $\varepsilon = \frac{1}{11m}$, the claim follows. □

Proof of Lemma 3.9. Let $k = \arg\min_{i\in T} \|\hat{a}_i\|_Q$. Let us use the decomposition of R as in Lemma 3.5. Then

$$\|\hat{a}_k\|_Q^2 \sum_{i=1}^n \gamma_i \le \sum_{i=1}^n \gamma_i \|\hat{a}_i\|_Q^2 = \sum_{i=1}^n \gamma_i \left(a_i^\mathsf{T} Q a_i/\|a_i\|_2^2\right) = \mathrm{tr}\left(Q \sum_{i=1}^n \gamma_i a_i a_i^\mathsf{T}/\|a_i\|_2^2\right) =$$
$$\mathrm{tr}(Q(R - \alpha I_m)) = \mathrm{tr}(I_m - \alpha Q) = m - \alpha \mathrm{tr}(Q) < m. \quad (14)$$

The second equality used Lemma 1.1(i) and (ii), the third used the decomposition of R, the fourth used $QR = I_m$, and the final inequality is since Q is positive definite.

The fact that $\mathrm{tr}(R) = \alpha m + \sum_{i=1}^n \gamma_i \le m + \sum_{i=1}^n \gamma_i$ and Lemma 1.1(iv) imply that $\sum_{i=1}^n \gamma_i \ge \mathrm{tr}(R) - m \ge m(\det(R)^{1/m} - 1)$. Note that the latter term is positive because $\det(R) > 1$, therefore the statement follows from (14). □

3.3 Maximum support case

The Full Support Image Algorithm naturally extends to the maximum support case. We assume a lower-bound $\bar{\omega}$ is given on $\hat{\omega}_A$. Lemma 3.7 easily generalizes to show $\|\hat{a}_k\|_Q \ge \hat{\omega}_A$ for every $k \in T^*$. Hence if the algorithm does not terminate within $O(m \log \bar{\omega}^{-1})$ rescalings, then we can identify a column $\|\hat{a}_k\|_Q < \bar{\omega}$ and conclude $k \notin T^*$. We thus eliminate a_k and recurse. A naïve approach would be to restart the algorithm every time a column is removed, but this would multiply the running time by a factor m. In what follows, we show that the same running time bound applies for the maximum support

Algorithm 6 Maximum Support Image Algorithm

Input: A matrix $A \in \mathbb{R}^{m\times n}$ with $\mathrm{rk}(A) = m$, and a lower-bound $\bar{\omega}$ on the condition number $\hat{\omega}_A$.

Output: A solution $\bar{y} \in \mathbb{R}^m$ to (I) satisfying the maximum number of strict inequalities.

Set $Q := I_m$, $R := I_m$, $U := I_m$, $T := [n]$, $r := m$.
while $T \neq \emptyset$ **do**
  Call VON NEUMANN(A, Q, ε) to obtain (x, y).
  if $A^\mathsf{T}Qy > 0$ **then return** $\bar{y} = UQy$. **Terminate.**
  else rescale

$$R := \frac{1}{1+\varepsilon}\left(R + \sum_{i\in T} \frac{x_i}{\|a_i\|_Q^2} a_i a_i^\mathsf{T}\right); \quad Q := R^{-1}.$$

  while $\exists k \in T$ such that $(\|\hat{a}_k\|_Q < \bar{\omega})$ or $(y = 0$ and $x_k > 0)$ **do**
    Select $W \in \mathbb{R}^{r\times(r-1)}$ whose columns form an orthonormal basis of $a_k^\perp$.
    Set $A := W^\mathsf{T}A$, delete all 0 columns, and remove the corresponding indices from T.
    Set $R := W^\mathsf{T}RW$, $U := UW$, and $r := r - 1$. Recompute $Q = R^{-1}$.
return $\bar{y} = 0$.

case as for the full support case. When a column is removed, we continue with the appropriate projection of the system, and keep track of the potential $\det(R)$ throughout the algorithm. The key new argument (Lemma 3.14) gives a bound on the potential decrease at any column removal.

Similarly to the Maximum Support Kernel Algorithm (Algorithm 2), we maintain a set T of indices with the property $T^* \subseteq T$. The set T is initialized as $T = [n]$, and we remove an index a_k once we conclude that $k \notin T^*$. The algorithm terminates with a solution $\bar{y}$ such that $a_k^\mathsf{T}\bar{y} > 0$ for all $i \in T$ and $a_k^\mathsf{T}\bar{y} = 0$ for all $i \notin T$, verifying $T = T^*$ at termination. We maintain r as the number of rows of A throughout the algorithm. As in the full support case, we assume that initially the matrix has full row rank; this will be preserved throughout the reduction steps.

Removing a column a_k from T is slightly more complicated than in the primal setting. We need to make sure $a_k^\mathsf{T}\bar{y} = 0$ for the desired solution $\bar{y}$, that is, we recurse on the subspace $a_k^\perp$. To do this we apply an isometric transformation of $a_k^\perp$ to $\mathbb{R}^{r-1}$. This is achieved by selecting an $r \times (r-1)$ matrix W whose columns form an orthonormal basis of $a_k^\perp$, and replacing A by the matrix obtained by removing all zero columns from $W^\mathsf{T}A$ (in particular, the kth column is removed since $W^\mathsf{T}a_k = 0$); thereby the number of rows decreases by one.

Theorem 3.11. *Let the matrix $A \in \mathbb{R}^{m\times n}$ have $\mathrm{rk}(A) = m$, and we are given a lower-bound $\bar{\omega}$ on $\hat{\omega}_A$. Algorithm 6 finds a maximum support solution to $A^\mathsf{T}y \geq 0$ in $O\left(m^2n^2\log(n\bar{\omega}^{-1})\right)$ arithmetic operations. The smoothed variant uses $O\left(m^3n\sqrt{\log n}\cdot\log(n\bar{\omega}^{-1})\right)$ arithmetic operations. If $A \in \mathbb{Z}^{m\times n}$ is integer of encoding size L, the running times of the two variants can be bounded as $O\left(m^2n^2L\right)$ and $O\left(m^3n\sqrt{\log n}\cdot L\right)$, respectively.*

3.3.1 Analysis

Note that the function $y \mapsto W^\mathsf{T}y$ is an isometric bijection from $\mathbb{R}^r \cap a_k^\perp$ to $\mathbb{R}^{r-1}$ whose inverse is the function $y' \mapsto Wy'$. The corresponding update to the matrix R will be

$W^\mathsf{T} RW$. Note that $E_{a_k^\perp}(R) = WE(W^\mathsf{T} RW)$ and that $\Sigma_A = W\Sigma_{W^\mathsf{T} A}$.

The matrix U keeps track of the product of all rescalings so far. We let $\bar{A}$ denote the original input matrix and A the current matrix at any iteration of the algorithm. The current matrix A is obtained from $U^\top \bar{A}$ by removing all zero columns (which are the columns outside T). Letting $H := \{y \in \mathbb{R}^m : \bar{a}_i^\top y = 0 \quad \forall i \in [n] \setminus T\}$, the transformation $y \mapsto U^\top y$ is an isometric bijection from H to $\mathbb{R}^r$ whose inverse is the function $y' \mapsto Uy'$. It is easy to see that $\Sigma_{\bar{A}} = U\Sigma_A$ and $F_{\bar{A}} = UF_A$.

We now formulate the main lemmas used in the proof of Theorem 3.11. The proof of the lemmas are deferred after the proof of the theorem. The first main lemma guarantees the correctness of the column removal.

Lemma 3.12. *If $F_A \subseteq E(R)$, then $\|\hat{a}_i\|_Q \geq \hat{\omega}_A$ for every $i \in T^*$. If $y = \sum_{\in T} x_i \hat{a}_i = 0$, $x \geq 0$, then $x_k \notin T^*$ for all k with $x_k > 0$.*

The next lemma is a strengthened version of Lemma 3.5, with explicit bounds on the coefficients. Note that the dimension m is replaced by the actual dimension r and the set of columns $[n]$ by T. Recall that $\bar{\omega}$ denotes a known lower bound on $\hat{\omega}_A$ for the initial input matrix A.

Lemma 3.13. *At any stage of the algorithm, we can write the matrix R in the form*

$$R = \alpha I_r + \sum_{i \in T} \gamma_i \hat{a}_i \hat{a}_i^\mathsf{T}$$

where $\gamma_i \leq 2/\bar{\omega}^2$, $\forall i \in T$, $\alpha = 1/(1+\varepsilon)^t$ for the total number of rescalings t performed thus far, and $\gamma_i \geq 0$. The trace is $\operatorname{tr}(R) = \alpha r + \sum_{i \in T} \gamma_i$. Furthermore, the matrix Q satisfies that for any unit-normed vector $v \in \mathbb{R}^r$, $\|v\|_Q \geq \bar{\omega}/\sqrt{2(n+1)}$.

The next lemma is the key new element in the argument for the maximum support problem.

Lemma 3.14. *Assume that at a given iteration $F_A \subseteq E(R)$, and consider an index $k \in T \setminus T^*$. Let $W \in \mathbb{R}^{r \times (r-1)}$ be a matrix whose columns form an orthonormal basis of $a_k^\perp$. Let A' be the matrix obtained by removing all zero columns from $W^\mathsf{T} A$, and let $R' = W^\mathsf{T} RW$. Then R' is positive definite and $F_{A'} \subseteq E(R')$. Furthermore,*

$$\det(R') \geq \frac{\bar{\omega}^2}{2(n+1)} \det(R).$$

We now prove Theorem 3.11 based on these lemmas and the results proved in Section 3.2.1.

Proof of Theorem 3.11. Lemmas 3.6 and 3.8 remain valid. Lemma 3.9 also holds under the assumption $\det(R) > 1$; the proof uses Lemma 3.13 in place of Lemma 3.5

We first show that the solution $\bar{y}$ returned by the algorithm is a solution to (I) satisfying the maximum number of strict inequalities. Lemmas 3.6 and 3.12 imply that $T \subseteq T^*$ throughout the algorithm. If $T = \emptyset$ at termination, then $\bar{y} = 0$ is indeed a maximum support solution. Assume the algorithm terminated at line 5 with $\bar{y} = UQy$, where the matrix $U \in \mathbb{R}^{m \times r}$ represents the sequence of transformations. For the original matrix $\bar{A}$, the current matrix A is obtained from $U^\mathsf{T} \bar{A}$ by removing all columns outside T; further, if $k \notin T$ then $U^\mathsf{T} \bar{a}_k = 0$. It follows that $\bar{a}_k^\mathsf{T} \bar{y} = 0$ for all such columns, and $\bar{a}_k^\top \bar{y} > 0$ for all $k \in T$, as required.

The potential $\det(R)$ is initially 1, and increases at least by a factor $16/9$ at every rescaling by Lemma 3.8. By Lemmas 3.9 and 3.12, we can find an index $k \in T \setminus T^*$ whenever $\det(R) > (1+\bar{\omega}^{-2})^m$. Furthermore, by Lemma 3.14, $\det(R)$ drops by at most a factor $\frac{\bar{\omega}^2}{2(n+1)}$ after the elimination of a column. Since the number of rows decreases by 1 every time we eliminate a column, the algorithm performs at most m column eliminations.

Consequently, within $O(m \log(n\bar{\omega}^{-1}))$ rescalings, all columns outside T^* will be removed and the algorithm terminates with a solution of maximum support.

As in the proof of Theorem 3.2, the iterations between two rescalings can be implemented in time $O(n^2 m)$, giving the running time bound. Following the proof of Theorem 3.4 we obtain the running time bound for the Smoothed Perceptron or MPFP updates. The bound on integer matrices follows by Lemma 1.4. □

The rest of this section is dedicated to the proof of the lemmas.

Proof of Lemma 3.12. The second statement is straightforward; we only prove the first part. We denote the original matrix and its columns as $\bar{A} = (\bar{a}_1, \ldots, \bar{a}_n)$. The same argument as in the proof of Lemma 3.7 shows that

$$\|\bar{a}_i\|_Q / \|\bar{a}_i\| \geq \hat{\omega}_A \quad \forall i \in T^*. \tag{15}$$

We recall that, for every $i \in T$, $a_i = U^\mathsf{T} \bar{a}_i$ for the current matrix U in the algorithm, and that U^T is an isometry. Thus $\|a_i\|_Q = \|\bar{a}_i\|_Q$. In what follows, we show that $\|a_i\| \leq \|\bar{a}_i\|$. Thus the claim follows from (15), using that $\|\hat{a}_i\|_Q = \|a_i\|_Q / \|a_i\| \geq \|\bar{a}_i\|_Q / \|\bar{a}_i\|$.

To show $\|a_i\| \leq \|\bar{a}_i\|$, we let $H := \{y \in \mathbb{R}^m : \bar{a}_i^\mathsf{T} y = 0 \quad \forall i \in [n] \setminus T\}$. Then $H \supseteq \Sigma_{\bar{A}}$ holds. We note that the matrix $\Pi := UU^T$ is the orthogonal projection matrix onto H. In particular, for all $i \in T$,

$$\|a_i\| = \sqrt{\bar{a}_i \Pi \bar{a}_i} \leq \sqrt{\|\bar{a}_i\| \|\Pi \bar{a}_i\|} \leq \|\bar{a}_i\|,$$

where the first inequality follows from Cauchy-Schwartz and the second from $\|\Pi \bar{a}_i\| \leq \|\bar{a}_i\|$. □

Proof of Lemma 3.13. The proof is by induction. The formula and bound is valid at initialization when $R = I_m$ and $\gamma_i = 0 \; \forall i \in [n]$. Let $R = \alpha I_r + \sum_{i \in T} \gamma_i \hat{a}_i \hat{a}_i^\mathsf{T}$ denote the current decomposition, where $\gamma_i \leq 2/\bar{\omega}^2$. We will show that the required form and bounds hold for the next update.

Assume that we rescale in the current iteration. Given this, we must have that $\min_{i \in T} \|\hat{a}_i\|_Q \geq \bar{\omega}$. Next, for $i \in [n]$, using Lemma 2.15 we see that

$$\|\hat{a}_i\|_Q^2 = \mathrm{width}^2_{E(R)}(\hat{a}_i) = \max \left\{ (\hat{a}_i^\mathsf{T} x)^2 : \alpha \|x\|^2 + \sum_{j \in T} \gamma_j (\hat{a}_j^\mathsf{T} x)^2 \leq 1, x \in \mathbb{R}^r \right\} \leq \frac{1}{\gamma_i}.$$

Now let x be the convex combination returned by Von Neumann in line 3. By the rescaling formula in line 5, the matrix R is updated to R' satisfying

$$R' = \frac{1}{1+\epsilon} \left(R + \sum_{i \in T} \frac{x_i}{\|a_i\|_Q^2} a_i a_i^\mathsf{T} \right) = \frac{1}{1+\epsilon} \left(\alpha I_r + \sum_{i \in T} \left(\gamma_i + \frac{x_i}{\|\hat{a}_i\|_Q^2} \right) \hat{a}_i \hat{a}_i^\mathsf{T} \right).$$

Therefore, each γ_i is updated to γ_i' satisfying

$$\gamma_i' = \frac{1}{1+\epsilon} \left(\gamma_i + \frac{x_i}{\|\hat{a}_i\|_Q^2} \right) \leq \frac{2}{\|\hat{a}_i\|_Q^2} \leq \frac{2}{\bar{\omega}^2},$$

as needed.

Consider now a step when some columns are eliminated. Then the matrices A and R are updated to A' and R', where A' is obtained by removing the zero columns from $W^\top A$ and $R' = W^\mathsf{T} R W$. We denote by $T' \subseteq T$ the index set of columns of A'. Thus

$$R' = \alpha W^\mathsf{T} W + \sum_{i \in T'} \gamma_i W^\mathsf{T} \hat{a}_i \hat{a}_i^\mathsf{T} W = \alpha I_{r-1} + \sum_{i \in T'} \gamma_i \|W^\mathsf{T} \hat{a}_i\|^2 \frac{a_i' a_i'^\mathsf{T}}{\|W^\mathsf{T} \hat{a}_i\|^2},$$

where the last equality follows from $W^\mathsf{T}W = I_{r-1}$ and the fact that $W^\mathsf{T}a_i = 0$ for all $i \in T \setminus T'$. Setting $\alpha' = \alpha$ and $\gamma'_i = \gamma_i \|W^\mathsf{T}\hat{a}_i\|^2$ for $i \in T'$ gives the desired decomposition of R'. Next, since $\|W^\mathsf{T}\hat{a}_i\| \leq \|\hat{a}_i\| \leq 1$, we get that $\gamma'_i \leq \gamma_i \leq 2/\bar{\omega}^2$, for all $i \in T'$.

We now prove the last part lower bounding $\|v\|_Q$ for any unit vector $v \in \mathbb{R}^r$. Firstly, for any $x \in \mathbb{R}^r$, the Cauchy-Schwarz inequality gives

$$x^\mathsf{T}Rx = \alpha\|x\|^2 + \sum_{i\in T} \gamma_i (\hat{a}_i^\top x)^2 \leq \left(\alpha + \sum_{i\in T}\gamma_i\right)\|x\|^2 \ ,$$

and hence $E(R)$ contains a Euclidean ball of radius at least $1/\sqrt{\alpha + \sum_{i\in T}\gamma_i}$. Therefore, for any unit vector $v \in \mathbb{R}^r$, using Lemma 2.15 we get

$$\begin{aligned}\|v\|_Q &= \max\left\{v^\top x : x \in E(R)\right\} \\ &\geq \frac{1}{\sqrt{\alpha + \sum_{i\in T}\gamma_i}} \geq \frac{1}{\sqrt{1 + 2|T|/\bar{\omega}^2}} \geq \frac{\bar{\omega}}{\sqrt{2(n+1)}} \ ,\end{aligned}$$

as needed. □

Proof of Lemma 3.14. The positive definiteness of R' follows easily from Lemma 3.5. To show that $F_{A'} \subseteq E(R')$, consider $y' \in F_{A'}$, and let $y = Wy'$. Observe that, by construction, $A^\mathsf{T}y = {A'}^\top y'$, $\|y\|_2 = \|y'\|_2$, and $y^\mathsf{T}Ry = (y')^\mathsf{T}R'y'$. This implies that $A^\mathsf{T}y \geq 0$ and $\|y\|_2 \leq 1$ because $y' \in F_{A'}$. It follows that $y \in F_A$, thus $y^\mathsf{T}Ry \leq 1$ because $F_A \subseteq E(R)$. This implies that $(y')^\mathsf{T}R'y' \leq 1$, meaning $y' \in E(R')$, as required.

Finally, note that $\det(R') = \det_{a_k^\perp}(R)$. Lemma 2.17 gives $\det(R') = \det(R)\|\hat{a}_k\|_Q^2$. To obtain the desired bound, we use the estimate $\|\hat{a}_k\|_Q^2 \geq \bar{\omega}^2/(2(n+1))$ from Lemma 3.13, which holds since $\hat{a}_k$ is a unit vector. □

4 Condition measures and orthonormal preconditioning

In this section we show how the condition measures we introduced in Section 1.2 relate to other conditions measure considered in the literature. We also address the question of finding a good preconditioning for our algorithms. For both the kernel and image problems, we get an equivalent problem, if we replace A by $B = TA$ for some non-singular $T \in \mathbb{R}^{m\times m}$. Hence the questions is whether we can replace A by another B with $\mathrm{im}(B^\top) = \mathrm{im}(A^\top)$, so that the relevant condition numbers improve. We show that, if B is chosen as a matrix whose rows form an orthonormal basis of $\mathrm{im}(A^\mathsf{T})$, many of our condition measures become approximately the best possible, and all running times can be lower bounded by Stewart's condition measure $\bar{\chi}_A$.

4.1 Relations to other condition measures

As already indicated in Section 1.2, our measures θ_A, ρ_A^*, and $\hat{\omega}_A$ are closely related to well known condition measures in the literature. The following pair of condition measures introduced by Vavasis and Ye [36]:

$$\begin{aligned}\sigma_A^K &\stackrel{\text{def}}{=} \min_{i\in S^*} \max\left\{x_i \,:\, x \in \ker(A)_+, \|x\|_1 = 1\right\}, \\ \sigma_A^I &\stackrel{\text{def}}{=} \min_{i\in T^*} \max\left\{x_i \,:\, x \in \mathrm{im}(A^\top)_+, \|x\|_1 = 1\right\}.\end{aligned} \tag{16}$$

Observe that the above measures depend only on the subspace $\mathrm{im}(A^\mathsf{T})$, rather than on the specific choice of the matrix A whose rows span the space. The relationship between θ_A and σ_A^K, and between $\hat{\omega}_A$ and σ_A^I, is formulated in the next claim,

Claim 4.1. *For every matrix $A \in \mathbb{R}^{m\times n}$, $\theta_A = \sigma_A^K/(1-\sigma_A^K)$. Furthermore, if $\|a_i\| = 1$ for all $i \in [n]$, then $\sigma_A^I \geq \hat{\omega}_A/n$.*

Proof. The first part of the claim is essentially the same as Proposition 22 in [16]. For the second part, under the assumption that $\|a_i\| = 1$ for all $i \in [n]$, we have $\|A^\top y\|_1 \leq n\|y\|$ for every $y \in \mathbb{R}^m$. The statement now from the definition of $\hat{\omega}_A$ and by observing that $\sigma_A^I = \min_{i\in T^*} \max\{a_i^\top y \ : \ y \in \mathbb{R}^m, A^\top y \geq 0, \|A^\mathsf{T} y\|_1 \leq 1\}$, and that $\hat{a}_i = a_i$ for all $i \in [n]$. □

The next lemma, shows the relationship between θ_A and ρ_A.

Claim 4.2. *Let $A \in \mathbb{R}^{m\times n}$ with $l = \max_{i\in[n]} \|a_i\| > 0$. Then, $\theta_A \geq -\rho_A/l$.*

Proof. Clearly, the statement is trivial if $\rho_A \geq 0$, thus we may assume $\rho_A < 0$. In this case, by Lemma 1.2 $\mathrm{conv}(A)$ contains a Euclidean ball of radius $|\rho_A|$. Thus for each $i \in [n]$, we have that $-(|\rho_A|/l)a_i \in |\rho_A|\mathbb{B}^m \subseteq \mathrm{conv}(A)$. Hence $\theta_A \geq -\rho_A/l$ as needed. □

Let us now introduce the condition numbers χ_A and $\bar{\chi}_A$ defined by Stewart [35] and by O'Leary [27]. Denoting by $\mathcal{D}$ the set of $n \times n$ diagonal matrices with positive entries, we define

$$\chi_A \stackrel{\text{def}}{=} \sup\left\{ \frac{\|y\|}{\|c\|} \ : \ y \text{ minimizes } \|D(A^\top y - c)\| \text{ for some } c \in \mathbb{R}^n, D \in \mathcal{D} \right\}$$
$$\bar{\chi}_A \stackrel{\text{def}}{=} \sup\left\{ \frac{\|A^\top y\|}{\|c\|} \ : \ y \text{ minimizes } \|D(A^\top y - c)\| \text{ for some } c \in \mathbb{R}^n, D \in \mathcal{D} \right\},$$

Note that the condition number $\bar{\chi}_A$ depends only on the subspace $\mathrm{im}(A^\top)$, whereas χ_A also depends on the matrix. In what follows, we focus on χ_A, and show that it is essentially equivalent to ρ_A^*. The $\bar{\chi}$ measure plays an important role in evaluating the running time of interior point methods (see Cheung et al. [6]), as well as in Section 4.2, where we show that it provides a lower bound on all of our condition measures, after an orthonormal preconditioning.

Let us characterize χ_A in terms of singular values. For a matrix B, let $s_{\min}(B)$ denote the smallest nonzero singular value of B.

Theorem 4.3 ([35, 27])**.** $1/\chi_A = \min_{\emptyset\neq S\subseteq[n]} s_{\min}(A_S)$.

Claim 4.4. *For every matrix $A \in \mathbb{R}^{m\times n}$, $1/(\sqrt{n}\chi_A) \leq \rho_A^* \leq 1/\chi_A$.*

Proof. Recall that $\mathbb{B}^d$ denotes the d-dimensional Euclidean ball. Let $\mathbb{B}_1^d = \{x \in \mathbb{R}^d : \|x\|_1 \leq 1\}$ denote the d-dimensional ℓ_1-ball. Consider a subset $S \subseteq [n]$ with $|S| = d$. Note that $\mathrm{conv}(A_S, -A_S) = A_S\mathbb{B}_1^d$, that is, the image of $\mathbb{B}_1^d$ under the linear transformation A_S. Lemma 1.2 gives

$$|\rho(A_S, -A_S)| = \max\left\{ r \ : \ r\mathbb{B}^m \cap \mathrm{im}(A_S) \subseteq A_S\mathbb{B}_1^d \right\}.$$

On the other hand, we can write $s_{\min}(A_S)$ as

$$s_{\min}(A_S) = \max\left\{ r \ : \ r\mathbb{B}^m \cap \mathrm{im}(A_S) \subseteq A_S\mathbb{B}^d \right\},$$

where $A_S\mathbb{B}^d$ is the image of $\mathbb{B}^d$ under A_S. Using that $\mathbb{B}_1^d \subset \mathbb{B}^d \subset \sqrt{n}\mathbb{B}_1^d$, we thus have $|\rho(A_S, -A_S)| \leq s_{\min}(A_S) \leq \sqrt{n}|\rho(A_S, -A_S)|$. The statement follows using the previous theorem. □

4.2 Orthonormal preconditioning

Given $A \in \mathbb{R}^{m\times n}$, throughout this section we consider a matrix $U = [u_1, \dots, u_n]$ whose rows form an orthonormal basis of $\mathrm{im}(A^\mathsf{T})$, that is, $\mathrm{im}(A^\mathsf{T}) = \mathrm{im}(U^\mathsf{T})$ and $UU^\mathsf{T} = I$. Observe that $\Pi_A^I = U^\mathsf{T}U$, therefore $\|u_i\| \le 1$ for all $i \in [n]$, since $\|u_i\|$ is the length of the projection of the ith unit vector onto $\mathrm{im}(A^\mathsf{T})$. A simple concrete way of obtaining an orthonormal basis is setting $U := (AA^\top)^{-1/2}A$.

We will be interested in evaluating how the condition measures we introduced change if we adopt the matrix U instead of A to represent our kernel and image feasibility problems. First, note that, in light of Claim 4.1, the measure θ_A depends only on the subspace $\mathrm{im}(A^\mathsf{T})$, rather than the specific choice of matrix A, hence θ_A does not change. On the other hand, we can expect improvement for the measures ρ_A^* and $\hat\omega_A$.

We show that the choice of U is "essentially optimal" with respect to the measures ρ and ρ^*, in the following sense. Let us assume that $\|a_i\| \le 1$ for all $i \in [n]$, and let $\mathcal{A}$ the set of all matrices $B = [b_1, \dots, b_n]$ such that $\mathrm{im}(B^\mathsf{T}) = \mathrm{im}(A^\mathsf{T})$ and $\|b_i\| \le 1$ for all $i \in [n]$. Then

$$\rho_U^* \ge \frac{1}{n}\sup_{B\in\mathcal{A}} \rho_B^*, \qquad |\rho_U| \ge \frac{1}{\sqrt{n}}\sup_{B\in\mathcal{A}} |\rho_B|. \tag{17}$$

The first bound is proved in Lemma 4.5. The second bound is trivial if $\rho_A = 0$; the nontrivial cases $S^* = [n]$ or $T^* = [n]$ follow by Lemma 4.8. Whereas we cannot prove the analogous statement for $\hat\omega_A$, we show that $\hat\omega_U$ can be lower bounded by σ_A^I (Claim 4.6).

We first observe that χ_U is independent from the choice of U, since $\|U^\top y\| = \|y\|$ for an orthonormal basis; moreover, $\chi_U = \bar\chi_A$.

Lemma 4.5. *If $\|a_i\| \le 1$ for all $i \in [n]$, then $\bar\chi_A \le \sqrt{n}\chi_A$ and $\rho_U^* \ge \rho_A^*/n$.*

Proof. Consider $x \in \mathrm{im}(A^\mathsf{T})$, $c \in \mathbb{R}^n$, and $D \in \mathcal{D}$ such that $A^\top y$ minimizes $\|D(A^\top y - c)\|$ and $\bar\chi_A = \|A^\top y\|/\|c\|$. Then

$$\bar\chi_A = \frac{\|A^\mathsf{T} y\|}{\|c\|} \le \frac{\|A^\mathsf{T}\|\|y\|}{\|c\|} \le \sqrt{n}\chi_A;$$

where $\|A^\mathsf{T}\|$ denotes the operator norm of A^T, and the last inequality follows from the definition of χ_A and the fact that $\|A\| \le \sqrt{n}$ because $\|a_i\| \le 1$ for all $i \in [n]$. Finally, from Claim 4.4, we have $\rho_U^* \ge 1/(\sqrt{n}\bar\chi_A) \ge 1/(n\chi_A) \ge \rho_A^*/n$. ☐

In particular, the above statement proves the first of the two inequalities in (17). The next claim shows that $\hat\omega_U$ is independent of the choice of the orthonormal basis U, and compares $\hat\omega_U$ to σ_A^I.

Claim 4.6. *We have*

$$\hat\omega_U = \min_{i\in T^*} \max\left\{ \frac{x_i}{\|u_i\|} \,:\, x \in \mathrm{im}(A^\top)_+, \|x\| = 1 \right\}.$$

Consequently, $\sigma_A^I \le \hat\omega_U$. Furthermore, if $\|a_i\| = 1$ for all $i \in [n]$, then also $\hat\omega_U/n \le \sigma_A^I$.

Proof. Since $\|U^\mathsf{T} y\| = \|y\|$ for all $y \in \mathbb{R}^m$, it follows that $\hat\omega_U = \min_{i\in T^*} \max\{x_i/\|u_i\| : x \in \mathrm{im}(A^\mathsf{T})_+, \|x\| \le 1\}$. The bound $\sigma_A^I \le \hat\omega_U$ follows from the definition of σ_A^I, using $\|u_i\| \le 1$, and that $\|x\| \le \|x\|_1$.

For the last part, we first show that if $\|a_i\| = 1\ \forall i \in [n]$, then $\|u_i\| \ge \frac{1}{\sqrt{n}}$. For $i \in [n]$, note that

$$\max_{x\in\mathrm{im}(A^\mathsf{T}),\|x\|=1} x_i = \max_{y\in\mathbb{R}^m,\|y\|=1} {u_i}^\top y = \|u_i\| \ .$$

where the first equality follows since rows of U form an orthornormal basis of $\mathrm{im}(A^\mathsf{T})$. Using that each a_i has norm 1, we also have that

$$\max_{x\in\mathrm{im}(A^\mathsf{T}),\|x\|=1} x_i \ge \frac{{a_i}^\top a_i}{\|A^\mathsf{T} a_i\|} \ge \frac{1}{\sqrt{n}},$$

as needed. The last part now follows combining the fact that $\|u_i\| \geq \frac{1}{\sqrt{n}}$, $\forall i \in [n]$, and that $\|x\|_1 \leq \sqrt{n}\|x\|$ $\forall x \in \mathbb{R}^n$. □

Hence, while it is not possible to lower-bound on $\hat{\omega}_A$ only in terms of σ_A^I and n, choosing an orthonormal basis U guarantees that $\hat{\omega}_U$ is no worse than σ_A^I.

Based on the previous bounds, we can bound the running time of all our algorithms in terms of n, m, and the single condition number $\bar{\chi}_A$, if an orthonormal preconditioning is applied.

Theorem 4.7. *Let A be a matrix and U be a matrix whose rows form an orthonormal basis for $\mathrm{im}(A^\mathsf{T})$. If we apply Algorithm 2 to U, the algorithm returns a maximum-support point in $\ker(A)_+$ in $O\left((m^3n + mn^2)\log(n\bar{\chi}_A)\right)$ arithmetic operations. If we apply Algorithm 6 to U, the algorithm returns a maximum-support point in $\mathrm{im}(A^\mathsf{T})_+$ in $O\left(m^2n^2\log(\bar{\chi}_A)\right)$ arithmetic operations.*

Proof. Vavasis and Ye [36, Theorem 5] showed that $\sigma_A^K, \sigma_A^I \geq 1/(\bar{\chi}_A + 1)$. Hence we have $\theta_U \geq 1/(\bar{\chi}_A + 1)$ (by Claim 4.1), $\rho_U^* \geq 1/(\sqrt{n}\bar{\chi}_A)$ (by Claim 4.4), and $\hat{\omega}_U \geq 1/(\bar{\chi}_A + 1)$ (by Claim 4.6). The bounds now follow by Theorems 2.8 and 3.11. □

The measure ρ_A and the full support kernel algorithm. We now focus on ρ_A. We assume that either $S^* = [n]$ or $T^* = [n]$, as otherwise $\rho_A = 0$. For the full support kernel case $S^* = [n]$, we show that orthonormal preconditioning leads to important conceptual improvements of the algorithm. In particular, we can bound the running time in terms of the larger condition number $|\hat{\rho}_U|$. The next lemma shows that, again, ρ_U is independent of the choice of the orthonormal U, and that choosing an orthonormal matrix is optimal up to a factor $\sqrt{n}$ with respect to this condition measure, verifying the second part of (17).

Lemma 4.8. *We have*

$$\rho_U = \max_{x \in \mathrm{im}(A^\mathsf{T})\setminus\{0\}} \min_{i\in[n]} \frac{x_i}{\|x\|}.$$

Furthermore, if $\|a_i\| \leq 1$ for all $i \in [n]$, then $|\rho_U| \geq |\rho_A|/\sqrt{n}$.

Proof. The first statement follows immediately from the definition of ρ_U and the fact that $\mathrm{im}(A^\mathsf{T}) = \mathrm{im}(U^\mathsf{T})$ and that $\|y\| = \|U^\mathsf{T}y\|$ for ever $y \in \mathbb{R}^m$.

For the second statement, assume first that $T^* = [n]$, and hence that both $\rho_A, \rho_U > 0$. Now consider $\bar{y} \in \mathrm{im}(A)$ such that $\|\bar{y}\| = 1$ and $\rho_A = \min_{i\in[n]} a_i^\mathsf{T}\bar{y}$, and let $\bar{x} = A^\mathsf{T}\bar{y}$. Note that $\rho_A = \min_{i\in[n]} \bar{x}_i > 0$ and that $\|\bar{x}\| \leq \sqrt{n}$ since $\|a_i\| \leq 1$ for all $i \in [n]$. Thus,

$$\rho_U \geq \min_{i\in[n]} \bar{x}_i/\|\bar{x}\| \geq \rho_A/\sqrt{n} > 0 \ .$$

Now assume that $S^* = [n]$ and hence $\rho_U, \rho_A < 0$. Now choose $\bar{x} \in \mathrm{im}(A^\mathsf{T})$, $\|\bar{x}\| = 1$, such that $\rho_U = \min_{i\in[n]} \bar{x}_i$. By definition, there exists $\bar{y} \in \mathrm{im}(A)$ s.t. $\bar{x} = A^\mathsf{T}\bar{y}$. Again as above, $1 = \|\bar{x}\| \leq \sqrt{n}\|\bar{y}\|$. Thus, we have that

$$0 > \rho_A \geq \min_{i\in[n]} \bar{x}_i/\|\bar{y}\| \geq \sqrt{n} \min_{i\in[n]} \bar{x}_i = \sqrt{n}\rho_U \ .$$

Taking absolute values, we get $|\rho_A| \leq \sqrt{n}|\rho_U|$ as needed. □

Next we focus on the relationship between ρ_U and σ_A^I, σ_A^K, in the cases $S^* = [n]$ or $T^* = [n]$.

Lemma 4.9.

- *If* $T^* = [n]$*, then*
$$\frac{1}{\sqrt{n}}\rho_U \leq \sigma_A^I \leq n\rho_U$$
- *If* $S^* = [n]$*, then*
$$\frac{|\rho_U|}{|\rho_U|+1} \leq \sigma_A^K \leq n|\rho_U| \ .$$

Proof. **(i).** The proof follows the same lines as the proof of Claim 1.3. The direction $\rho_U \leq \sqrt{n}\sigma_A^I$ is straightforward by Lemma 4.8 and the definition of σ_A^I, and using that $\|x\|_1 \leq \sqrt{n}\|x\|$. For the other direction, let $x^{(i)} := \arg\max\{x_i : x \in \mathrm{im}(A^\top)_+, \|x\|_1 = 1\}$ for all $i \in [n]$; hence $\sigma_A^I = \min_{i\in[n]} x_i^{(i)}$. We let $\bar{x} = \sum_{i=1}^n x^{(i)}/n$; clearly, $\|\bar{x}\|_1 = 1$, and therefore $\|\bar{x}\| \leq \|\bar{x}\|_1 = 1$. Thus $\rho_U \geq \sigma_A^I/n$.

(ii). We first prove $\sigma_A^K \leq n|\rho_U|$. The statement will follow from the following stronger bound.

Claim 4.10. *Assume* $S^* = [n]$*. For every* $z \in \ker(A)_+$ *with* $\|z\|_1 = 1$*, and every* $x \in \mathrm{im}(A^\mathsf{T})$ *with* $\|x\|_2 = 1$*, there exists* $k \in [n]$ *such that* $x_k \geq \min_{i\in[n]} z_i$*.*

Proof. Let $\tau = \min_{i\in[n]} z_i$ and suppose by contradiction that $x_i < \tau$ for all $i \in [n]$. Since $T^* = \emptyset$, it follows that x has at least one negative component. Then $\sum_{i\,:\,x_i>0} x_i z_i < (1-\tau)\tau$. On the other hand, since $\sum_{i=1}^n x_i z_i = 0$,

$$\sum_{i\,:\,x_i>0} x_i z_i = \sum_{i\,:\,x_i<0} |x_i| z_i \geq \tau \sqrt{\sum_{i\,:\,x_i<0} x_i^2} = \tau\sqrt{1 - \sum_{i\,:\,x_i>0} x_i^2} > \tau\sqrt{1-(n-1)\tau^2}.$$

It follows that $1-(n-1)\tau^2 < (1-\tau)^2$, which immediately implies $\tau > 2/n$, a contradiction since $\tau \leq 1/n$. ☐

To show $|\rho_U| \geq \sigma_A^K/n$, let $z^{(i)} := \arg\max\{z_i : z \in \ker(A)_+, \|z\|_1 = 1\}$ for every $i \in [n]$, so that $\sigma_A^K = \min_{i\in[n]} z_i^{(i)}$. Let $z = \frac{1}{n}\sum_{i=1}^n z^{(i)}$, so that $z \in \ker(A)_+$, $\|z\|_1 = 1$, and $z_i \geq \sigma_A^K/n$ for all $i \in [n]$. Now for any $x \in \mathrm{im}(A^\mathsf{T})$, $\|x\| = 1$, applying Lemma 4.8 to $-x$ we must have that $\min_{i\in[n]} x_i \leq -\min_{i\in[n]} z_i \leq -\sigma_A^K/n < 0$. It then follows that $|\rho_U| \geq \sigma_A^K/n$.

We now show that $\sigma_A^K \geq \frac{|\rho_U|}{|\rho_U|+1}$. Firstly, by Claim 4.1, $\sigma_A^K = \frac{\theta_A}{\theta_A+1}$. Next, we recall that $\theta_A = \theta_U$ since $\ker(A) = \ker(U)$. Lastly, by Claim 4.2, noting that $\rho_U < 0$ and $\max_{i\in[n]} \|u_i\| \leq 1$, we have that $\theta_U \geq |\rho_U|$ and hence $\frac{\theta_U}{\theta_U+1} \geq \frac{|\rho_U|}{|\rho_U|+1}$. The claim thus follows. ☐

Finally, we show that, for the case $S^* = [n]$, if we apply Algorithm 1 with U as an input matrix, then the algorithm will terminate with a positive solution as soon as the vector y computed during the execution has norm less than 1. This is a substantial improvement over the original analysis, where we can only conclude this once $\|y\| < |\rho_{(A,-A)}|$.

Claim 4.11. *Given* $x \in \mathbb{R}^n$ *such that* $x \geq \vec{e}$*, let* $y = U^\mathsf{T}x$*. If* $\|y\| < 1$*, then* $\pi_A^K x > 0$*.*

Proof. Since $UU^\mathsf{T} = I$, it follows that $\Pi_A^I = U^\mathsf{T}U$, thus $\|\Pi_A^I x\| = \|y\| < 1$. Since $x \geq \vec{e}$, it follows that $\Pi_A^K x = x - \Pi_A^I x > 0$. ☐

Recall from the analysis of Algorithm 1 that the algorithm performs at most $m\log_{3/2}|\hat{\rho}_U|^{-1}$ when applied to matrix U. Since, by the previous claim, the algorithm terminates as soon as the vector y computed during the execution has norm less than 1, it follows that the total number of DV updates depends on $\hat{\rho}_U$, rather than ρ_U. We summarize this in the following corollary.

Corollary 4.12. *Let A be a matrix such that (K_{++}) is feasible, and U be a matrix whose rows form an orthonormal basis for $\mathrm{im}(A^{\mathsf{T}})$. Algorithm 1 applied to matrix U finds a feasible solution of (K_{++}) in $O(m^2 n \log n + m^3 n \log |\hat{\rho}_U|^{-1})$ arithmetic operations.*

Recalling that $\hat{\rho}_U \geq \rho_U$ because all columns of U have norm at most 1, and recalling that $\rho_U \geq \rho_A/n$, it follows that the above running time is not worse than that given by any choice of input matrix A. Although this does not give an improvement in terms of the worst-case running time, it might be beneficial in practice.

References

[1] S. Agmon. The relaxation method for linear inequalities. *Canadian Journal of Mathematics*, 6(3):382–392, 1954.

[2] A. Basu, J. A. D. Loera, and M. Junod. On Chubanov's method for linear programming. *INFORMS Journal on Computing*, 26(2):336–350, 2013.

[3] A. Belloni, R. M. Freund, and S. Vempala. An efficient rescaled perceptron algorithm for conic systems. *Mathematics of Operations Research*, 34(3):621–641, 2009.

[4] U. Betke. Relaxation, new combinatorial and polynomial algorithms for the linear feasibility problem. *Discrete & Computational Geometry*, 32(3):317–338, 2004.

[5] L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin (New Series) of the American Mathematical Society*, 21(1):1–46, 1989.

[6] D. Cheung, F. Cucker, and Y. Ye. Linear programming and condition numbers under the real number computation model. *Handbook of Numerical Analysis*, 11:141–207, 2003.

[7] S. Chubanov. A polynomial relaxation-type algorithm for linear programming. http://www.optimization-online.org/DB_FILE/2011/02/2915.pdf, 2011.

[8] S. Chubanov. A strongly polynomial algorithm for linear systems having a binary solution. *Mathematical programming*, 134(2):533–570, 2012.

[9] S. Chubanov. A polynomial algorithm for linear optimization which is strongly polynomial under certain conditions on optimal solutions. http://www.optimization-online.org/DB_FILE/2014/12/4710.pdf, 2015.

[10] S. Chubanov. A polynomial projection algorithm for linear feasibility problems. *Mathematical Programming*, 153(2):687–713, 2015.

[11] D. Dadush, L. A. Végh, and G. Zambelli. Rescaled coordinate descent methods for linear programming. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 26–37. Springer, 2016.

[12] G. B. Dantzig. Converting a converging algorithm into a polynomially bounded algorithm. Technical report, Stanford University, 1991.

[13] G. B. Dantzig. An ε-precise feasible solution to a linear program with a convexity constraint in $1/\varepsilon^2$ iterations independent of problem size. Technical report, Technical Report 92-5, Stanford University, 1992.

[14] J. Dunagan and S. Vempala. A simple polynomial-time rescaling algorithm for solving linear programs. *Mathematical Programming*, 114(1):101–114, 2008.

[15] M. Epelman and R. M. Freund. Condition number complexity of an elementary algorithm for computing a reliable solution of a conic linear system. *Mathematical Programming*, 88(3):451–485, 2000.

[16] M. Epelman and R. M. Freund. A new condition measure, preconditioners, and relations between different measures of conditioning for conic linear systems. *SIAM Journal on Optimization*, 12(3):627–655, 2002.

[17] M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.

[18] J. Goffin. The relaxation method for solving systems of linear inequalities. *Mathematics of Operations Research*, 5(3):388–414, 1980.

[19] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.

[20] R. Hoberg and T. Rothvoß. An improved deterministic rescaling for linear programming algorithms, 2016. manuscript.

[21] D. Li, C. Roos, and T. Terlaky. A polynomial column-wise rescaling von Neumann algorithm. http://www.optimization-online.org/DB_FILE/2015/06/4979.pdf, 2015.

[22] D. Li and T. Terlaky. The duality between the perceptron algorithm and the von neumann algorithm. In *Modeling and Optimization: Theory and Applications*, pages 113–136. Springer, 2013.

[23] T. Motzkin and I. Schoenberg. The relaxation method for linear inequalities. *Canadian Journal of Mathematics*, 6(3):393–404, 1954.

[24] A. Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

[25] Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.

[26] A. B. Novikoff. On convergence proofs for perceptrons. In *Proceedings of the Symposium on the Mathematical Theory of Autoomata XII*, pages 615–622, 1962.

[27] D. P. O'Leary. On bounds for scaled projections and pseudoinverses. *Linear Algebra and its Applications*, 132:115–117, 1990.

[28] J. Peña, D. Rodriguez, and N. Soheili. On the von Neumann and Frank–Wolfe algorithms with away steps. *SIAM Journal on Optimization*, 26(1):499–512, 2016.

[29] J. Peña and N. Soheili. Solving conic systems via projection and rescaling. *arXiv preprint arXiv:1512.06154*, 2015.

[30] J. Peña and N. Soheili. A deterministic rescaled perceptron algorithm. *Mathematical Programming*, 155(1-2):497–510, 2016.

[31] K. Roos. On Chubanov's method for solving a homogeneous inequality system. In *Numerical Analysis and Optimization*, pages 319–338. Springer, 2015.

[32] F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386–408, 1958.

[33] A. Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, New York, 1998.

[34] N. Soheili and J. Peña. A smooth perceptron algorithm. *SIAM Journal on Optimization*, 22(2):728–737, 2012.

[35] G. W. Stewart. On scaled projections and pseudoinverses. *Linear Algebra and its Applications*, 112:189–193, 1989.

[36] S. A. Vavasis and Y. Ye. Condition numbers for polyhedra with real number data. *Operations Research Letters*, 17(5):209–214, 1995.

[37] L. A. Végh and G. Zambelli. A polynomial projection-type algorithm for linear programming. *Operations Research Letters*, 42(1):91–96, 2014.

[38] P. Wolfe. Finding the nearest point in a polytope. *Mathematical Programming*, 11(1):128–149, 1976.

[39] A. W. Yu, F. Kılınç-Karzan, and J. G. Carbonell. Saddle points and accelerated perceptron algorithms. In *Proceedings of The 31st International Conference on Machine Learning – Journal of Machine Learning Research 3*, pages 1827–1835, 2014.

A Appendix

Lemma 1.2. *$|\rho_A|$ equals the distance of 0 from the relative boundary of* $\mathrm{conv}(A)$*. Further,*

(i) $\rho_A < 0$ if and only if 0 is in the relative interior of $\mathrm{conv}(A)$*, or equivalently, $S^* = [n]$.*

(ii) $\rho_A > 0$ if and only if 0 is outside $\mathrm{conv}(A)$*, or equivalently, $T^* = [n]$. In this case, the Goffin measure $\hat{\rho}_A$ equals the* width of the image cone Σ_A*, that is, the radius of the largest ball in $\mathbb{R}^m$ centered on the surface of the unit sphere and inscribed in Σ_A.*

Proof. By Lemma A.1 below, $|\rho_A|$ is the distance of 0 from the relative boundary of $\mathrm{conv}(A)$. **(i)** By Lemma A.1(ii), $\rho_A < 0$ if and only if 0 is in the relative interior of $\mathrm{conv}(A)$, which is the case if and only if there exists $x > 0$ such that $Ax = 0$.

(ii) For any $\bar{y} \in \Sigma_A$, $\|\bar{y}\| = 1$, the distance between $\bar{y}$ and the hyperplane $\{y : a_j^\mathsf{T} y = 0\}$ ($j \in [n]$) is $\hat{a}_j^\mathsf{T} \bar{y}$, therefore $\min_{j\in[n]} \hat{a}_j^\mathsf{T} \bar{y}$ is the distance of $\bar{y}$ from the boundary of Σ_A, that is, the radius of the largest ball centered at $\bar{y}$ and contained in Σ_A. The statement now follows from the definition of $\hat{\rho}_A$. □

Lemma A.1. *Let $A \in \mathbb{R}^{m\times n}$. Let p be a point of minimum norm in the relative boundary of* $\mathrm{conv}(A)$*.*

(i) If $0 \notin \mathrm{conv}(A)$, then $\|p\| = \rho_A = \min_{j\in[n]} a_j^\mathsf{T} \hat{p}$.

(ii) If 0 is in the relative interior of $\mathrm{conv}(A)$*, then p is in the relative interior of some facet of* $\mathrm{conv}(A)$ *and $\|p\| = -\rho_A = \max_{j\in[n]} a_j^\mathsf{T} \hat{p}$.*

Proof. **(i)** Assume $0 \notin \mathrm{conv}(A)$. Then p is a point of minimum norm in $\mathrm{conv}(A)$. It follows that $p^\mathsf{T} z \geq \|p\|^2$ for every $z \in \mathrm{conv}(A)$, implying that $\|p\| \leq \min_{j\in[n]} a_j^\mathsf{T} \hat{p} \leq \rho_A$. We now show that $\rho_A \leq \|p\|$. If not, then there exists $y \in \mathrm{im}(A)$ such that $\|y\| = 1$ and $\min_{j\in[n]} a_j^\mathsf{T} y > \|p\|$. In particular, this implies that every point in $\mathrm{conv}(A)$ has distance greater than $\|p\|$ from the origin, contradicting our choice of $p \in \mathrm{conv}(A)$.

(ii) Assume 0 is in the relative boundary of $\mathrm{conv}(A)$. By our choice of p, for any $y \in \mathrm{im}(A)$, $\|y\| = 1$, we have $\|p\|y \in \mathrm{conv}(A)$, thus there exist $x \geq 0$ such that $\vec{e}^\mathsf{T} x = 1$ and $\|p\|y = Ax$. It follows that $\|p\| = \sum_{j=1}^n x_j a_j^\mathsf{T} y$, therefore there exists $j \in [n]$ such that $a_j^\mathsf{T} y \geq \|p\|$. This shows $-\rho_A \geq \|p\|$.

Consider now the minimal face F of $\mathrm{conv}(A)$ containing p, and let $J = \{j \in [n] : a_j \in F\}$. Then $p = \sum_{j\in J} a_j \lambda_j$ where $\lambda_j > 0$, and $\sum_{j\in J} \lambda_j = 1$. Thus every $y \in F$ satisfies $p^\mathsf{T} y = \|p\|^2$. We argue that F is a facet of $\mathrm{conv}(A)$. Suppose not. Then there exists a facet F' strictly containing F. Furthermore, there exists some $j \in [n]$ such that $p^\mathsf{T} a_j < \|p\|^2$ because $0 \in \mathrm{conv}(A)$, therefore we may choose F' such that there exists $k \in [n]$ with $a_k \in F'$ and $a_k^\mathsf{T} p < \|p\|^2$. Consider the point $p' = \lambda a_k + (1-\lambda)p$ where $\lambda = p^\mathsf{T}(p - a_k)/\|p - a_k\|^2$. Note that $\lambda > 0$ because $p^\mathsf{T} a_k < \|p\|^2$, furthermore, since $\|a_k\| \geq \|p\|$ (otherwise a_k would be a point on the relative boundary of $\mathrm{conv}(A)$ with smaller norm) we also have $p^\mathsf{T} a_k \leq \|a_k\|^2$, which implies $\lambda < 1$. It follows that p' is a convex combination of two points in F', therefore $p' \in F'$. Moreover, $\|p'\|^2 = \|p\|^2 - \lambda^2\|p - a_k\|^2 < \|p\|^2$, therefore p' is a point on the relative boundary of $\mathrm{conv}(A)$ with smaller norm, a contradiction.

This implies that p is in the relative interior of a facet F of $\mathrm{conv}(A)$, therefore $p^\mathsf{T} y \leq \|p\|^2$ is a valid inequality for $\mathrm{conv}(A)$ and $F = \mathrm{conv}(A) \cap \{y : p^\mathsf{T} y = \|p^2\|\}$. This implies that $a_j^\mathsf{T} \hat{p} \leq \|p\|$ for all $j \in [n]$. On the other hand, F contains at least one element a_k, $k \in [n]$, therefore $a_k^\mathsf{T} \hat{p} = \|p\|$. It follows that $\|p\| = \max_{j\in[n]} a_j^\mathsf{T} \hat{p}$. □

For the purpose of proving Lemma 1.4, when estimating the bit complexity of a matrix A with integer entries and encoding size L, we will refer to the quantity

$$\Delta_A = \max_B \prod_{j \in B} \|a_j\|,$$

where $B \subseteq [n]$ ranges over all sets such that $|B| = \mathrm{rk}(A)$ and the columns of A_B are linearly independent. It can be easily shown that $\Delta_A < 2^L$ (see for example [19, Lemma 1.3.3]). By Hadamard's bound, for every square submatrix A' of A, $|\det(A')| \leq \Delta_A$.

Lemma 1.4. *Assume $A \in \mathbb{Z}^{m \times n}$ of total encoding length L. Then $\theta_A, \rho_A^*, \hat{\omega}_A \geq 2^{-O(L)}$. If $\rho_A \neq 0$, then $|\rho_A|, |\hat{\rho}_A| \geq 2^{-O(L)}$.*

Proof. Let S^*, T^* be defined as in (1).

Claim A.2. *If $T^* \neq \emptyset$, then $\max_{y \in \Sigma_A \setminus \{0\}} \min_{j \in T^*} a_j^\mathsf{T} \hat{y} \geq \dfrac{1}{m\Delta_A}$.*

Let $(y^*, s^*) \in \mathbb{R}^{2m}$ be an optimal solution of the following linear program:

$$\begin{array}{lrcl} \min & \vec{e}^\mathsf{T} s & & \\ & A_{T^*}^\mathsf{T} y & \geq & \vec{e} \\ & A_{S^*}^\mathsf{T} y & = & 0 \\ & s - y & \geq & 0 \\ & s + y & \geq & 0 \end{array}$$

This LP is feasible by the definition of T^*; the optimal value equals $\|y^*\|_1$. The vector y^* is a basic solution of the system $A_{T^*}^\mathsf{T} y \geq \vec{e}$, $A_{S^*}^\mathsf{T} y = 0$, therefore $|y_j^*| \leq \Delta_A$ for all $j \in [n]$, so $\|y^*\|_1 \leq \sqrt{m}\Delta_A$.

Since by construction $y^* \in H \setminus \{0\}$, the statement follows from the fact that

$$\min_{j \in T^*} a_j^\mathsf{T} \frac{y^*}{\|y^*\|_2} = \frac{1}{\|y^*\|_2} \geq \frac{1}{\sqrt{m}\|y^*\|_1} \geq \frac{1}{m\Delta_A}.$$

This concludes the proof of the claim.

Observe that, if we let $\alpha := \max_{i \in T^*} \|a_i\|$, we have

$$\hat{\omega}_A = \min_{j \in T^*} \max_{y \in \Sigma_A \setminus \{0\}} \hat{a}_j^\mathsf{T} \hat{y} \geq \alpha^{-1} \min_{j \in T^*} \max_{y \in \Sigma_A \setminus \{0\}} a_j^\mathsf{T} \hat{y} \geq \alpha^{-1} \max_{y \in \Sigma_A \setminus \{0\}} \min_{j \in T^*} a_j^\mathsf{T} \hat{y} \geq \frac{1}{m\Delta_A^2},$$

where the last inequality follows from $\alpha \leq \Delta_A$ and Claim A.2. It follows that $\hat{\omega}_A \geq 2^{-O(L)}$.

Claim A.3. *If $\rho_A \neq 0$, then $|\rho_A| \geq \dfrac{1}{m^m \Delta_A}$.*

If $0 \notin \mathrm{conv}(A)$, then Claim A.2 is applicable with $S^* = \emptyset$ and $H = \mathbb{R}^m$, hence we get the stronger $\rho_A \geq 1/(m\Delta_A)$. Assume now that 0 is in the relative interior of $\mathrm{conv}(A)$. Let $r := \mathrm{rk}(A)$. By $\rho_A \neq 0$, 0 must be in the relative interior of $\mathrm{conv}(A)$ and therefore Lemma A.1(ii) is applicable. Consider a point p of minimum norm in the relative boundary of $\mathrm{conv}(A)$. Then $|\rho_A| = \|p\|$, and p is in the relative interior of a facet F of $\mathrm{conv}(A)$, where F is defined by the valid inequality $p^\mathsf{T} y \leq \|p\|^2$. Thus, given an $m \times r$ submatrix B of A whose columns are linearly independent elements of F, it follows that p is the point of minimum norm in the hyperplane H of $\mathrm{im}(A)$ affinely generated by B, that is, $H = \{y : y = Bz,\ \exists z \in \mathbb{R}^r,\ \vec{e}^\mathsf{T} z = 1\}$. This shows that $p = \|p^2\| B(B^\mathsf{T} B)^{-1} \vec{e}$, therefore $\|p\|^2 = (\vec{e}^\mathsf{T} (B^\mathsf{T} B)^{-1} \vec{e})^{-1}$.

Let $M := B^\mathsf{T} B$. Given $P, Q \subseteq [r]$, we denote by $M_{P,Q}$ the submatrix of M defined by the rows indexed by P and by the columns indexed by Q. By the Cauchy-Binet formula, for every $1 \leq t \leq r$ and every choice of $P, Q \subseteq [r]$ such that $|P| = |Q| = t$

$$\det(M_{P,Q})^2 = \sum_{\substack{U \subseteq [m] \\ |U| = t}} \det(B_{P,U}) \det(B_{Q,U}) \leq \binom{m}{t} \Delta_A^2 \leq m^t \Delta_A^2.$$

By Cramer's rule, for every $i, j \in [r]$, $M^{-1}_{ij} = (-1)^{i+j}\dfrac{\det(M_{[r]\setminus\{i\},[r]\setminus\{j\}})}{\det(M)}$. It follows that $|M^{-1}_{ij}| \leq |\det(M_{[r]\setminus\{i\},[r]\setminus\{j\}})| \leq m^{r/2}\Delta_A$, therefore

$$\rho_A = \frac{1}{\sqrt{\vec{e}^{\mathsf{T}} M^{-1}\vec{e}}} \geq \frac{1}{\sqrt{r^2\sqrt{m^{r/2}}\Delta_A}} \geq \frac{1}{m^m\sqrt{\Delta_A}}.$$

This concludes the proof of the claim.

If follows from Claim A.3 that $|\rho_A| \geq 2^{-O(L)}$ if $\rho_A \neq 0$. Also, since $|\hat{\rho}_A| \geq |\rho_A|/(\max_{j\in[n]} \|a_j\|)$ and $\max_{j\in[n]} \|a_j\| \leq \Delta_A$, we have $|\hat{\rho}_A| \geq 2^{-O(L)}$ if $\rho_A \neq 0$. Furthermore, note that, for every $S \subseteq [n]$, $\Delta_{(A_S,-A_S)} \leq \Delta_A$, therefore Claim A.3 also implies that $\rho^*_A \geq 1/(m^m\Delta_A)$, and thus $\rho^*_A \geq 2^{-O(L)}$.

To bound θ_A, observe that $\mathrm{conv}(A) \cap \mathrm{im}(A_{S^*}) \supseteq \mathrm{conv}(A_{S^*})$, and by Lemma 1.2 $\mathrm{conv}(A_{S^*})$ contains the ball $(|\rho_{A_{S^*}}|)\mathbb{B}^m \cap \mathrm{im}(A_{S^*})$. It then follows from the definition of θ_A that $\theta_A \geq (\min_{i\in S^*} \|a_i\|)|\rho_{A_{S^*}}|$, thus $\theta_A \geq 1/(m^m\Delta^2_{A_{S^*}}) \geq 1/(m^m\Delta^2_A)$ if $S^* \neq \emptyset$. In particular, $\theta_A \geq 2^{-O(L)}$ if $S^* \neq \emptyset$. □

Lemma 2.1. *Let $A \in \mathbb{R}^{m\times n}$ and $S^* = S^*_A$. Then $\mathrm{span}(P_A) = \mathrm{im}(A_{S^*})$, and $P_A = P_{A_{S^*}}$.*

Proof. We first show $P_A = P_{A_{S^*}}$. The inclusion $P_{A_{S^*}} \subseteq P_A$ is obvious. For the reverse inclusion, consider $y \in P_A$ and let $x, z \in \mathbb{R}^n_+$ such that $\vec{e}^{\mathsf{T}}x = \vec{e}^{\mathsf{T}}z = 1$ and $y = \hat{A}x = -\hat{A}z$. Then $\hat{A}(x + z) = 0$, $x + z \geq 0$, which implies $x_i = z_i = 0$ for all $i \in [n] \setminus S^*$, which shows that $y \in P_{A_{S^*}}$.

We show $\mathrm{span}(P_A) = \mathrm{im}(A_{S^*})$. It suffices to show that $\mathrm{span}(P_{A_{S^*}}) = \mathrm{im}(A_{S^*})$ because $P_A = P_{A_{S^*}}$. The inclusion $\mathrm{span}(P_{A_{S^*}}) \subseteq \mathrm{im}(A_S)$ is obvious. For the reverse inclusion, it suffices to show that, for every $i \in S$, there exists $\alpha \neq 0$ such that $\alpha a_i \in P_{A_{S^*}}$. Consider $\lambda \in \mathbb{R}^{|S^*|}_{++}$ such that $\hat{A}_{S^*}\lambda = 0$, and assume without loss of generality that $\sum_{j\in S^*\setminus\{i\}} \lambda_j = 1$. Then $-\lambda_i\hat{a}_i = \sum_{j\in S^*\setminus\{i\}} \lambda_j\hat{a}_j$, which implies $-\lambda_i\hat{a}_i \in P_{A_{S^*}}$. □

Lemma 2.5. *Let $X \in \mathbb{R}$ be a random variable supported on the interval $[-\varepsilon, \eta]$, where $0 \leq \varepsilon \leq \eta$, satisfying $\mathbb{E}[X] = \mu$. Then for $c \geq 0$, we have that*

$$\mathbb{E}[\sqrt{1 + cX^2}] \leq \sqrt{1 + c\eta(\varepsilon + |\mu|)}$$

Proof. Let $l(x) = \frac{\eta-x}{\eta+\varepsilon}\sqrt{1 + c\varepsilon^2} + \frac{x+\varepsilon}{\eta+\varepsilon}\sqrt{1 + c\eta^2}$ denote the unique affine interpolation of $\sqrt{1 + cx^2}$ through the points $\{-\varepsilon, \eta\}$. By convexity of $\sqrt{1 + cx^2}$, we have that $l(x) \geq \sqrt{1 + cx^2}$ for all $x \in [-\varepsilon, \eta]$. Hence, we see that

$$\begin{aligned}\mathbb{E}[\sqrt{1 + cX^2}] &\leq \mathbb{E}[l(X)] \quad (\text{since } X \text{ is supported on } [-\varepsilon, \eta])\\ &= l(\mathbb{E}[X]) = l(\mu) \quad (\text{since } l \text{ is affine}).\end{aligned}$$

From here, we get that

$$\begin{aligned}l(\mu) &= \frac{\eta-\mu}{\eta+\varepsilon}\sqrt{1 + c\varepsilon^2} + \frac{\mu+\varepsilon}{\eta+\varepsilon}\sqrt{1 + c\eta^2}\\ &\leq \sqrt{1 + c\left(\frac{\eta-\mu}{\eta+\varepsilon}\varepsilon^2 + \frac{\mu+\varepsilon}{\eta+\varepsilon}\eta^2\right)} \quad (\text{by concavity of } \sqrt{x})\\ &= \sqrt{1 + c\left(\eta\varepsilon + (\eta-\varepsilon)\mu\right)} \leq \sqrt{1 + c\eta(\varepsilon + |\mu|)} \quad (\text{since } \varepsilon \leq \eta),\end{aligned}$$

as needed. □